University of Groningen

Folk moral objectivism

Zijlstra, Lieuwe

DOI:
[10.33612/diss.109876462](https://doi.org/10.33612/diss.109876462)

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:
2019

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

Citation for published version (APA):
Zijlstra, L. (2019). *Folk moral objectivism*. [Thesis fully internal (DIV), University of Groningen]. University of Groningen. <https://doi.org/10.33612/diss.109876462>

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.

FOLK MORAL OBJECTIVISM

Lieuwe Zijlstra

Folk Moral Objectivism

PhD Thesis

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

to obtain the degree of PhD of Ghent University
on the authority of
Rector Prof. R. Van de Walle
and in accordance with
the decision by the College of Deans

Double PhD degree

This thesis will be defended in public on
Thursday 19 December 2019 at 09.00 hours
by

Lieuwe Zijlstra

born on 28 April 1983
in Hoogezand-Sappemeer, The Netherlands

Supervisors

Prof. F.A. Hindriks

Prof. J. Verplaetse

Co-supervisor

Dr. H.W.A. Evers

Assessment committee

prof. G. Björnsson

Prof. F. Mortier

Prof. B. Streumer

Prof. M.M.S.K. Sie

Dr. K. Reuter

Spread your wings and fly away

Fly away, far away

Pull yourself together

Cause you know you should do better

That's because you're a free man

Queen (written by John Deacon) - 1977

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	X
PREFACE	XI
INTRODUCTION	1
1. DOES EMPIRICAL RESEARCH DEBUNK FOLK MORAL OBJECTIVISM?	7
INTRODUCTION	8
1. FOLK MORAL OBJECTIVISM	9
2. PSYCHOLOGICAL RESEARCH ON FOLK METAETHICS	11
2.1. The very beginning of empirical research on folk metaethics	12
2.2. A novel experimental methodology: truth-aptness & disagreement	14
2.3 A different experimental design to obtain high construct validity	18
3. IMPLICIT AND EXPLICIT MORAL OBJECTIVISM	23
CONCLUSION	26
2. FOLK MORAL OBJECTIVISM AND ITS MEASUREMENT	29
1. INTRODUCTION	30
2. BEYOND PERCEIVED OBJECTIVITY	32
2.1. Objectivity as a multi-dimensional construct	32
2.2 The predictive power of folk moral objectivism	34
2.3 The psychological distinctiveness of moral attitudes	35
3. THE PRESENT RESEARCH	37
4. STUDY 1: DEVELOPMENT OF THE MEASUREMENT SCALE	38
4.1 Method	39
4.1.1 Participants	39
4.1.2 Materials and Procedure	39
4.1.3 Exploratory Factor Analysis	41
4.2 Discussion	43
5. STUDY 2: A CONFIRMATORY FACTOR ANALYSIS OF THE THREE-FACTOR STRUCTURE	45
5.1 Method	46
5.1.1 Participants	46
5.1.2 Materials and Procedure	46
5.1.3 Confirmatory Factor Analysis	46
5.2 Discussion	47

6. STUDY 3: FIRST VALIDATION STUDY	47
6.1 Method	48
6.1.1 Participants	48
6.1.2 Materials and Procedure	49
6.2 Results	49
6.3 Discussion	50
7. STUDY 4: SECOND VALIDATION STUDY	52
7.1. Methods	53
7.1.1 Participants	53
7.1.2 Materials and Procedure	53
7.2 Results	55
7.3 Discussion	60
8. STUDY 5: WILLINGNESS TO HARM	62
8.1 Methods	64
8.1.1 Participants	64
8.1.2 Materials and procedure	64
8.2 Results	66
8.3 Discussion	68
9. GENERAL DISCUSSION	69
9.1 Relationship to previous research on perceived objectivity	70
9.2 Relation to other research in social psychology	72
10. CONCLUSION	72
3. ARE PEOPLE IMPLICITLY MORAL OBJECTIVISTS?	75
1. ARE PEOPLE IMPLICITLY MORAL OBJECTIVISTS?	76
1.1 Enoch's Tests	78
2.1 Test 1: Jokes	79
2.2. Test 2: The phenomenology of disagreement	81
2.3. Test 3: The counterfactual question	82
3. EXPERIMENT	82
3.1. Methods	83
3.1.1. Participants	83
3.1.2. Materials and Procedure	84
3.1.3. Results	85
2.2. Discussion	88

4. CONCLUSION	89
APPENDIX 1: ATTENTION CHECKS	90
APPENDIX 2: RESULTS FIRST SURVEY EXPERIMENT	91
APPENDIX 3: THE PHENOMENOLOGY OF DISAGREEMENT	93
APPENDIX 4: THE COUNTERFACTUAL QUESTION	94
4. DO PEOPLE BELIEVE IN OBJECTIVE MORAL PROGRESS, KNOWLEDGE, AND ERROR?	97
1. INTRODUCTION	98
2. THE PRESENT RESEARCH	101
2.1 Method	102
2.1.1 Participants	102
2.1.2 Materials and procedure	102
2.1.3 Statements	105
2.3 Results	106
2.3.1 Statistical hypotheses and analytical strategy	106
2.3.2 Descriptives	107
2.3.3 Testing Hypothesis 1: Moral progress	109
2.3.4 Testing Hypothesis 2: Moral knowledge	110
2.3.5 Testing Hypothesis 3: Moral error	111
2.3.6 Hypothesis tests with subjective categorizations	112
3. DISCUSSION	113
4. CONCLUSION	114
CONCLUSION	117
APPENDIX A: HOW NOT TO ARGUE ABOUT IS/UGHT INFERENCES IN THE COGNITIVE SCIENCES	123
1. INTRODUCTION	123
2. KEY CONCEPTS	123
3. MIND THE GAP	125
4. A DEBATE SHACKLED	127
SAMENVAATING	131
CURRICULUM VITAE	139
BIBLIOGRAPHY	141

ACKNOWLEDGEMENTS

The research presented in this PhD thesis is conducted on the basis of a double degree contract between Ghent University and the University of Groningen. At Ghent University, this research was conducted as part of the FWO-project “Measuring Moral Relativism”. All of the survey experiments are funded by the Research Foundation Flanders (FWO) – project nr. G.0683.13N. As part of this project, a paper written by Katinka Quintelier and me was published in 2014 in *Frontiers in Psychology* titled “How (not) to argue about is/ought inferences in the cognitive sciences”. This paper can be found in Appendix A.

PREFACE

If I had to identify a central theme to describe the period in which I wrote this PhD thesis then it would be one about crossing borders and changing places. In a literal sense, all the work that I have done to complete this PhD thesis (i.e. running survey experiments, writing papers, meeting with advisors or co-authors, attending workshops, submitting papers to journals, and so forth) was done in Ghent, Groningen, Rotterdam, New Haven, Buffalo, New York, Graz, Bochum, and numerous other places. Parts of this thesis were even written in Kralendijk (Bonaire) and behind the steering wheel of a truck (of course, while parked in a loading dock!). During this time, I changed my home address seven times.

There is also a metaphorical sense in which this thesis crossed borders. The research that is part of this thesis is interdisciplinary in nature. I investigated an empirical claim that plays a role in philosophical theories and I employed research methodologies typically used by social scientists. In the process, I tried to translate philosophical assumptions into empirical hypotheses and I assessed the philosophical significance of empirical findings. In order to get this work done properly, I studied philosophical literature in metaphysics, psychological and sociological theories on moral values and behavior, and I consulted literature on statistical methodology to appropriately run the necessary statistical analyses. Moreover, I consulted with philosophers, sociologists, psychologists, statisticians and many others from a range of different disciplines.

This thesis could not have been completed, however, if there had not people who were willing to share their knowledge and skills, who advised me on how to write academic papers or how to run statistical analyses correctly, or who helped me organize and coordinate everything necessary for my stays abroad or for the double degree contract between Groningen and Ghent. One of the people who deserve a special mention is Joshua Knobe who so generously spent time tutoring me during my stay at Yale University. He provided

me with the opportunity to be a visiting assistant in research and he generously spent time to help me work out survey experiments, getting to know others in the research field, and advising me on how to write academic papers.

I am also very grateful for the support of my promotors and advisors, Jan Verplaetse and Katinka Quintelier in Ghent, Frank Hindriks, Daan Evers, and Bart Streumer in Groningen. They made it possible for me to work in both places and they advised me on all aspects of this PhD thesis. This thesis has improved considerably as a result of their comments, advice, recommendations, and by the discussions that we had about metaethics, experimental philosophy, and the relationship between both. I would also like to thank Thomas Pölzler and Jacob Dijkstra. It really was a pleasure to work together with them on developing research on moral progress, knowledge and error (Chapter 4) and I learned a lot from them.

With regard to all the coordination and organization of this project, I am very grateful for the support provided by Katelijne Verstichel and Marga Hids. There was a lot of work involved in arranging the different stays abroad, running experiments and arranging funding, and making it possible for this project to take place in Groningen, Ghent, and New Haven. Katelijne and Marga put a lot of time and effort in arranging these things and always did so in good spirit. I am grateful for that.

During my PhD research I also received a lot of support from my parents and I am obviously very grateful for that. The support that they provided made my life and work much easier, much more comfortable, and more pleasurable. I also have pleasant memories of their visits to each of the places that I spent a longer amount of time in (Groningen, Ghent, New Haven).

In the past years I have not only developed myself in an academic sense but it was also very much a personal journey. I am grateful for the friendship and support of Debora Weerman, Joringel den Hoedt, Kari Wheeler-Reed, Nick Stagnaro, and Ryan Mays. To Joringel, I would like to express my gratitude for his friendship (and all the extremely pleasant and fun moments in the past years) and for the fact that he has always been a true believer in what

I am capable of doing (at least from his perspective!). I share many interests with Joringel, especially those that make life so much interesting, which includes philosophy, music, literature, film, art, and comedy, and Joringel is the one who noted the relevance of the song *Spread your Wings* by Queen in relation to this thesis (see quote above). I am grateful to Debora for her joyous spirit and her patience and care in the past years and for the pleasant memories in Groningen, Ghent, New Haven, Bonaire, and many other places. To Kari, I would like to express my gratitude for supporting me while doing my PhD research, for the elaborate conversations that we had about the many meaningful and joyful things in life, and for her lasting friendship. I am extremely grateful for the fact that she tirelessly supported me in times when life was tough. Moreover, I will always be amazed by her exceptional abilities to put into words those things that are very difficult to perceive and/or to express verbally. I am also grateful to Nick Stagnaro and Ryan Mays for making me feel welcome in New Haven and for the many pleasurable moments that we experienced in New Haven and in Rotterdam/Amsterdam during their visits. Finally, I am grateful for the support of fellow PhD students in Groningen, Ghent, and Yale. I would especially like to thank Stipe Pandžić for being one of my paranymphs and for his friendship in the past years.

This journey has come to an end but I look forward to be en route toward something new where I can build on the work that I have done as part of this PhD.

INTRODUCTION

When it comes to the study of human morality, philosophers make a distinction between normative ethics, applied ethics, and metaethics. Theories in normative ethics focus on first-order questions about what actions are morally right or wrong and good or bad. Theories in applied ethics concern the practical application of ethical theories to real-life circumstances and scenarios. Theories in metaethics, on the other hand, aim at explaining what morality fundamentally is. For instance, they concern the content of moral concepts, the relationship between moral reasons and being motivated to act morally, whether there are moral properties and what kind of entity they are, and whether or not moral objectivism is true.

This PhD thesis presents an investigation into the question of whether an assumption often made by metaethicists is true, namely the thesis of “folk moral objectivism” (FMO). FMO is an empirical thesis about the nature of ordinary moral discourse and thought. According to this thesis, people believe that moral sentences are objectively true or false. To wit, that moral objectivism is true. More specifically, the thesis describes people as believing that moral sentences are true or false independently of any individual’s subjective attitudes and any particular cultural perspective. The arguments that philosophers provide for FMO often come in the form of descriptions of ordinary moral discourse and appeals to intuitions. Despite the fact that this is an empirical claim that can be tested, philosophers rarely use empirical methodology to test whether FMO is true.

The question of whether or not FMO is true is important for theories in metaethics. First, philosophers actively try to accommodate FMO in their metaethical theories. Second, some philosophers claim that the truth of FMO makes some metaethical theories more plausible (i.e. moral realism) and other metaethical theories less plausible (i.e. moral antirealism).

In recent years, psychologists and experimental philosophers have started to investigate FMO empirically (Nichols 2004; Goodwin & Darley 2008, 2010,

2012; Pölzler & Wright unpublished; Sarkissian, Park, Tien, Wright & Knobe 2011; Wright, Cullen & McWhine 2011). In extant studies, people treat many moral sentences as if those sentences are subjectively true or false or neither true nor false. Moreover, as I discuss extensively below, there are no participants in those studies who consistently treat moral sentences as objectively true or false. Given the large support that FMO receives from philosophers and their effort to accommodate FMO in their theories, these findings are remarkable and in need of explanation.

The main research question of this PhD thesis is, therefore, whether or not FMO is true. A subsidiary research question concerns the best way of measuring FMO. In this PhD manuscript you will find an investigation into both questions and tentative answers.

In Chapter 1, I review existing research on folk metaethics and I evaluate whether or not they provide evidence for or against FMO. My investigation reveals that it is not clear whether they provide genuine evidence. The first reason for this is that many measures do not directly or not exclusively measure what is relevant for FMO. The second reason is that there is a distinction between implicit and explicit metaethical commitments. For instance, people may explicitly believe, and verbally express, that morality is relative. At the same time, it is possible that they implicitly believe, and act accordingly, as if there are objective moral facts. Consequently, implicit and explicit metaethical commitments can diverge. I argue that some philosophers claim that FMO applies to implicit commitments and that it is not clear what type of commitment is measured by existing research on folk metaethics.

In Chapter 2 I present research in which I investigate whether an existing psychological construct, which is used in empirical research as an indicator of folk moral objectivism, measures moral objectivism versus non-objectivism or also related metaethical views. More specifically, I investigated whether people’s intuitions about moral truth are best measured on a single dimension of moral objectivism (i.e. perceived objectivity) versus non-objectivism or whether there are multiple psychological dimensions that underlie

people's intuitions. If there are different dimensions, is this related to people's scores on tolerance measures? For instance, although philosophically distinct views, it is psychologically possible that people's intuitions about the objectivity of moral judgments are on a single dimension with moral absolutism and opposed to relativism. If so, it may be true that higher scores on objectivism/absolutism decrease people's tolerance toward others and that lower scores (i.e., having intuitions about relativism) increase people's tolerance. However, it is also possible that, psychologically speaking, objectivism and absolutism are two distinct dimensions and that intuitions about relativism oppose the former and not the latter, or vice versa. If so, this also raises new questions about the relationship with tolerance.

Existing research has shown that there are large differences in people's judgments about the objectivity of moral judgments, both between different individuals and between different moral statements. My results strongly suggest that people's intuitions about moral truth are indeed multidimensional and that each dimension, which I have termed Independent Truth, Universal Truth, and Divine Truth, have different relationships with tolerance and a willingness to harm measure.

In Chapter 3 I argue that it is not clear whether measurement instruments used in existing research measure implicit or explicit commitments. I therefore present an investigation specifically focusing on people's implicit commitments. The reason for this investigation is twofold. First, FMO might be about implicit commitments, which some philosophers suggest is the case. Second, there is reason to think that previous research mostly measured explicit commitments. If this is true, the large differences that we observe in people's judgments about the truth and falsity of moral judgments might not be relevant for our assessment of FMO. Fortunately, Enoch (2014) has developed three tests, which he uses as intuition pumps, which in his view should show that people are implicitly committed to moral objectivism even if they explicitly deny this. I use those tests as material for a survey experiment to test whether people indeed respond in a way that Enoch expects. The results show

that for each of the thought experiments separately, the majority of people do respond as if FMO is true. However, when people's responses are combined, this becomes less clear. Nevertheless, this research provides some support for the idea that people are implicitly committed to moral objectivism.

To obtain more insight into people's metaethical commitments, I explore in Chapter 4 alternative ways of measuring FMO (in collaboration with Thomas Pölzler and Jacob Dijkstra). Previous research investigated whether or not people believe that moral judgments are objectively true or false and whether or not they believe that at most one party can be correct in a moral disagreement. If FMO is true, however, people will also believe in the possibility of moral progress, knowledge, and error (or so we argue). We therefore designed a survey experiment to investigate whether this is indeed the case. Our participants responded, both abstractly and for concrete cases, to questions about moral progress, knowledge, and error, in the domains of morality, science, personal preferences, and social conventions.

We assumed that people believe that scientific statements are objectively true or false. We also assumed that people do not believe that statements about personal preferences or social conventions are objectively true or false. Our results show that when people are asked abstractly, and also when asked to judge concrete cases, they provide responses that suggest that they believe scientific statements are objectively true or false. With regard to morality, however, people's responses are, overall, more comparable to their responses to personal preferences or conventions. Consequently, the results of this study do not provide evidence for FMO.

Overall, my research reveals that we should be careful in interpreting empirical results as evidence against FMO. First, previous studies that claim to provide support for folk moral objectivism do not always measure what is relevant for FMO. Second, the studies in this thesis do not provide univocal evidence for it. Third, an ultimate verdict requires a thorough investigation of the distinction between implicit and explicit commitments. This thesis takes the first step to measuring implicit commitments, again without providing

univocal support for FMO. Thus, in spite of initial optimism, I conclude that the jury is still out. If FMO is indeed true, this is not immediately obvious if we look at empirical data.

1.

**DOES EMPIRICAL RESEARCH
DEBUNK FOLK MORAL OBJECTIVISM?**

Abstract

A prevalent assumption among philosophers is that people are moral objectivists, which I refer to as the thesis of folk moral objectivism (FMO). Recent psychological research appears to shed doubt on FMO. For instance, even though people seem to treat some moral statements as objectively true or false, they treat other moral statements as subjectively true or false, or even as neither true nor false. Moreover, there seem to be large differences between different individuals with regard to the objectivity that they attribute to morality. In this paper I investigate whether the findings of psychological research provide evidence that shows that FMO is true or false. My analysis reveals, firstly, that many of the empirical measures that are used in psychological research do not directly pertain to FMO. Secondly, those measures that seem to pertain to FMO may elicit explicit commitments. According to some philosophers, however, FMO concerns people's implicit commitments. I conclude that the thesis of folk moral objectivism is not (yet) debunked by existing research.

Keywords: Folk moral objectivism, construct validity, implicit, explicit

Introduction

Some philosophers believe that there are objectively true moral judgments (Brink 1984; Gibbard 1992; Railton 1996; Smith 1994; Sturgeon 1985) and others do not (Ayer 1936; Hare 1954; Blackburn; 1984; Harman 1975; Strummer 2017). Despite the differences among philosophers about whether or not moral objectivism is true, most of them share the assumption that lay people believe in moral objectivism (Brink 1984; Brink 1989; Mackie 1977; Shaffer-Landau 2003; Smith 1994). Moreover, most philosophers believe that this is something that metaethical theories should accommodate. In this paper, I will refer to the assumption that ordinary people are committed to moral objectivism as 'the thesis of folk moral objectivism' (FMO).

Psychologists and experimental philosophers have recently started to empirically investigate folk metaethics and the results of those studies appear to make it less likely that FMO is true. Although people generally treat moral statements more objectively than judgments about matters of taste or social conventions (Goodwin and Darley 2008), they do not treat moral statements as objectively as factual statements, which are ordinarily considered as belonging to a domain that is characteristically objective (Goodwin and Darley 2008; Wright, Grandjean & McWhite 2013). More strikingly, there are large differences between different people in the objectivity they attribute to moral issues; and people treat certain moral issues objectively and other moral issues non-objectively (Goodwin & Darley 2008, 2010, 2012; Wright, McWhite & Grandjean 2013; Wright, Grandjean & McWhite 2013; Zijlstra, 2019).

There is, however, sufficient reason not to take these studies at face value. First, many measures that are used do not directly measure folk moral objectivity. Second, there is a difference between being explicitly or implicitly committed to moral objectivism (Zijlstra, unpublished). And it may be that only the latter type of commitment is relevant for FMO.

To show that this is the case, I will discuss some of the key findings of psychological research on folk metaethics and argue why they neither support nor undermine FMO. The paper is organized as follows. In section 1, I discuss

why philosophers believe that the FMO assumption is true. On the basis of these reasons, I develop three criteria of folk moral objectivity. In section 2, I discuss psychological research and I employ the three criteria in order to assess whether or not this research provides evidence in favour of FMO. I argue that it does not. In section 3, I argue that it is possible to be a moral objectivist in two distinct ways, namely implicitly and explicitly, and that existing psychological research has mostly measured explicit commitments. It is therefore still possible that people are moral objectivists, namely implicitly.

1. Folk Moral Objectivism

Many philosophers assume that people are moral objectivists and that this is something metaethical theories should accommodate (Brink, 1989; Cuneo, 2007; Finlay, 2007; Nichols, 2004). For example, Brink (1984) writes that “[I]n moral deliberation and argument we search for answers to moral questions, answers whose correctness we assume to be independent of our means of arriving at them” (Brink, 1989, p. 11). In his book, Brink also claims that people perceive moral claims as assertions that can be objectively true or false, that they believe that some people are better at grasping objective moral facts than others, and that they believe that objective moral progress is possible. In the rare case that people do respond as antirealists, this is due to a cognitive process that led them to believe that the realists’ commitments are untenable (Brink, 1989, p. 23).

This fits well with what Smith (1994) writes about how we (that is, ordinary people) think about morality:

We seem to think moral questions have correct answers, that the correct answers are made correct by objective moral facts; that moral facts are wholly determined by circumstances and that, by engaging in moral conversation and argument, we can discover what these objective moral facts determined by the circumstances are. (Smith, 1994, p. 6)

What is crucial is that these philosophers suggest that people think that moral sentences are true or false independently of whatever subjective responses they have. I will term this criterion ‘Independence’. Moreover, people are described as believing that those who have different moral views are mistaken. More specifically, that only one party to a moral disagreement can be correct (Jackson 2000, Smith 1994, Streumer 2018). I will call this criterion ‘Exclusion’.

An assumption that is usually left implicit is that, if one moral belief is objectively true or false, then other moral beliefs are also objectively true or false. Parfit questions this assumption:

We should not assume that the objectivity of Ethics must be all-or-nothing. There may be a part of morality that is objective. In describing this part, our claims may be true. When we consider this part of morality, or these moral questions, we may find the Unified Theory that would remove our disagreements. There may be other questions about which we shall never agree. There may be no true answers to these questions. Since objectivity need not be all-or-nothing, moral sceptics may be partly right. These questions may be subjective. (Parfit, 1984, p. 452)

Consequently, Parfit believes that the thesis of moral objectivism does not necessarily imply that all moral claims are objectively true or false. Instead, it is possible that some moral statements are objectively true or false and others are subjectively true or false. It is possible that the same applies for the commitments of the folk. People may believe that some moral claims are objectively true or false and others are not. Some philosophers and psychologists, however, believe that what we might call ‘the uniformity assumption’ is part of FMO (Goodwin & Darley 2008; Pölzler 2018; Pölzler & Wright, unpublished; Sinnott-Armstrong 2009). On this interpretation, a violation of the uniformity assumption provides evidence against FMO (Goodwin & Darley 2008; Pölzler 2018; Pölzler & Wright, unpublished; Sinnott-Armstrong 2009).

In my discussion of psychological research on folk metaethics in the next section I will use these three criteria to evaluate the evidence for or against FMO.

We have, then, the following core criteria for FMO:

Independence: People believe that moral sentences are true or false independently of anyone's subjective reactions or attitudes

Exclusion: People believe that only one party in a moral disagreement can be correct

Uniformity: Independence and Exclusion hold for all moral beliefs

Independence and Exclusion are both necessary for FMO. They are rejected by people who are moral subjectivists (moral sentences are true or false in virtue of the existence of subjective facts), by nihilists and error theorists (there are no true moral sentences), and by non-cognitivists (moral sentences are neither true nor false¹). In light of the fact that it is contested whether or not a violation of Uniformity is a violation of FMO, I will evaluate Uniformity in light of empirical research, but leave it to my reader to decide on its importance.

2. Psychological Research on Folk Metaethics

In my evaluation of FMO on the basis of Independence, Exclusion, and Uniformity, I will take specific empirical studies as point of departure. I will start with discussing research by Nichols (2004), which is the first empirical study that explicitly takes FMO as point of departure. Subsequently, I will discuss the results of research using truth-aptness and disagreement tasks, which have been developed by Goodwin and Darley (2008, 2012) and have become standard tests in this field of research. Finally, I will discuss research by Pölzler and Wright (unpublished) who used truth-aptness and disagreement tasks but also introduced additional measurement instruments to measure

¹ Some non-cognitivists endorse minimalist theories of truth. Minimalism does allow non-cognitivists to say that moral claims are truth-apt.

FMO. In the next section I will argue that despite a wealth of interesting results, there are two important reasons why we cannot yet interpret those results as providing clear evidence for or against FMO.

2.1. The very beginning of empirical research on folk metaethics

An early psychological study into the question of whether people are moral objectivists has been conducted by Nichols (2004). Nichols explored people's views of different kinds of disagreements (moral, conventional, factual). In the moral disagreement scenario, participants were presented with the following:

John and Fred are members of different cultures, and they are in an argument. John says "It's okay to hit people just because you feel like it," and Fred says "No, it is not okay to hit people just because you feel like it." John then says "Look you are wrong. Everyone I know agrees that it's okay to do that." Fred responds, "Oh no, you are the one who is mistaken. Everyone I know agrees that it's not okay to do that." (Nichols, 2004, p. 9)

Subsequently, participants were asked which of the following answer options best characterized their views:

- It is okay to hit people just because you feel like it, so John is right and Fred is wrong.
- It is not okay to hit people just because you feel like it, so Fred is right and John is wrong.
- There is no fact of the matter about unqualified claims like "It's okay to hit people just because you feel like it." Different cultures believe different things, and it is not absolutely true or false that it's okay to hit people just because you feel like it. (Nichols, 2004, p. 9-10)

Nichols assumes that selecting the first or second answer option reflects a commitment to moral objectivism. This is because selecting the first or sec-

ond answer option implies that there is a correct response to the moral disagreement. Moreover, Nichols presumably assumes that people who select the first or second answer option deny the third answer option. The third answer option contains the idea that there are no absolute facts that can arbitrate moral disagreements, which is a rejection of both Independence (because there are no objectively true or false moral sentences) and Exclusion (because both parties can be correct)². If Nichols' interpretation of the responses of his participants is correct, the first and second answer option measures Independence and Exclusion and the third answer denies both.

The results of his first study show that 17 out of 40 participants respond that there is no fact of the matter about hitting someone because you feel like it.³ In a second study, 9 out of 40 participants responded that there is no fact of the matter about the moral statement involving hitting people because you feel like it. In a different study, Nichols used a scenario involving aliens from a different planet who believe that it is ok to torture puppies for the fun of it. In this case, 15 out of 40 participants responded that there is no fact of the matter about the statements involving torturing puppies for fun. Consequently, between 23% and 43% of participants in Nichols' research respond in ways that are in tension with Independence and Exclusion, and between 57% and 77% respond in accordance with both criteria.

Consequently, some people seem to believe that there are no objective moral facts that can arbitrate moral disagreements and others seem to believe there are. If FMO is true, it is surprising to find that a sizable minority of people responds that there are no objective moral facts. At the same time, it is also true that a majority of people does respond in a way that corresponds to FMO. Of course, we should not base our assessment of FMO only on the research by Nichols (2004). In order to assess an empirical thesis like FMO,

² I assume here that non-philosophers interpret the term 'absolutely' in Nichols' study as being roughly similar to how I use the term 'objectively' in this paper.

³ Nichols also tested how participants respond to factual statements (i.e. whether the earth is flat) and excluded participants who believe that there is no fact of the matter about such statements. Strikingly, respectively 3, 6, and 12 participants for the first three studies responded that there is no fact of the matter about the flat-earth scenario.

we will need to evaluate a range of different studies and investigate it from different angles.

Moreover, the measurement instrument used by Nichols might be flawed in that participants interpret it in a way that has no bearing on FMO. For instance, it is not entirely clear whether or not people's responses to the first and second answer option measure what is relevant for FMO (i.e. namely whether or not they believe there are objectively true or false moral judgments) or instead merely reflects people's normative views (i.e. whether or not it is okay to hit someone because you feel like it). If it is the latter, the responses of participants provide no evidence for or against FMO.

Also, the moral example that Nichols (2004) uses as input for the scenario is fairly extraordinary. I conjecture that people will rarely, if ever, consider moral cases that concern the issue of whether or not it is okay to hit someone because they feel like it. The extraordinariness of this case may induce people to provide a response that does not reflect how they ordinarily respond to moral disagreements. Of course, one could also argue that such borderline cases are especially relevant for evaluating FMO. Indeed, philosophers often use quite radical examples to test their intuitions. In fact, the scenario used by Nichols (2004) contains direct and senseless harm, which makes it surprising that some people do not express their disapproval. Consequently, we need to evaluate additional research findings on folk metaethics before we can come to a conclusive assessment of whether or not FMO is true.

2.2. A novel experimental methodology: truth-aptness & disagreement

In this section I will start by discussing research by Goodwin and Darley (2008). Goodwin and Darley were the first to administer a truth-aptness and disagreement task, which *prima facie* seem to measure Independence and Exclusion. For the truth-aptness task, participants were presented with a range of statements from each of the domains and they were then asked whether those statements are true or false or whether they are opinions or attitudes instead. The options "True" and "False" are meant to measure Independence

while the option “opinions or attitudes” are interpreted as a rejection of Independence. For now, I will assume that this is correct. However, below I will return to this issue.

The results on the truth-aptness measure show that, on average, the six cases representing factual statements are each overwhelmingly rated as true or false (>90%). This is what we would expect from a domain that is characteristically objective (i.e. the domain of scientific facts). This is not what we observe for the moral cases. The nine different moral cases were on average rated only in 38% of the cases as true or false. Strikingly, this is almost similar to how participants treat conventional matters (45%). Unsurprisingly, matters of taste were not at all perceived as true or false (7%).

Most strikingly, people’s scores on the truth-aptness measure vary strongly on the basis of the content of the moral statement at issue. For example, statements about randomly shooting other people on the street or about cheating on an exam are respectively judged in 68% and 54% of the cases as true or false. Statements involving donating to charity (36%), assisting in the death of a terminally ill friend (8%), performing stem cell research (2%), and abortion (2%), are not at all judged as true or false but are perceived as opinion or attitudes instead.

In the second phase of the experiment participants were requested to respond to moral disagreements. Participants were presented with two statements they strongly agreed or disagreed with in the truth-aptness task. They were then told that someone else in the study disagreed strongly with them about the content of the statement. Subsequently, participants were asked to choose from the following options:

- (1) The other person is surely mistaken
- (2) It is possible that neither you nor the other person is mistaken
- (3) It could be that you are mistaken, and the other person is correct
- (4) Other

At first sight, this seems to be a straightforward measure of Exclusion and therefore of FMO. I will continue my evaluation as if this is indeed the case, but I will raise doubts about this assumption below.

Goodwin and Darley use a composite score of both the truth-aptness and disagreement task in their presentation of results. The composite score has three different categories: Fully objective, intermediately objective, and least objective. If participants perceived a moral statement to be true or false and regarded that at most one person can be correct in a moral disagreement, this was labelled as fully objective and applies to 50% of the cases. If participants rated a moral statement as true or false *or* regarded that at most one person can be correct in a disagreement about the statement, this was labelled as intermediately objective, which applies to 28% of the cases. If participants denied that a moral statement can be true or false and they allowed for the possibilities that both parties can be correct, this was labelled as the least objective response, which applies to 11% of the cases.⁴

What do Goodwin and Darley’s results imply for FMO? It seems that people do not consistently treat moral statements in a way that corresponds to Independence or moral disagreements in a way that corresponds to Exclusion. Hence, we observe that people do not treat moral statements in ways that correspond to these criteria. Moreover, we observe that there are large individual differences and that people’s responses strongly differ based on moral content. There are no individuals who consistently treat moral statements in a way that corresponds to Independence and Exclusion. There are also no moral statements that are consistently treated in a way that corresponds to both criteria. What we observe instead is that people treat some moral beliefs as true or false in accordance with Independence and Exclusion but not the whole set of moral beliefs. Hence, Uniformity does not apply to how people treat moral statements. The significance of this observation for FMO depends on whether you believe that the uniformity assumption is true or false. If you believe that the uniformity assumption is true then these findings appear to

⁴ In another 11% of the cases participants selected the answer option “Other” in the case of a moral disagreement, and those responses could therefore not be categorized.

show that people violate Uniformity and that FMO is false. If you believe that the uniformity assumption is false, it is possible to interpret these findings as support for FMO. That is, people appear to treat some moral statements in accordance with Independence and Exclusion and perhaps that is sufficient for FMO to be true.

However, does the truth-aptness task genuinely measure Independence and does the disagreement task measure Exclusion? As I wrote above, the truth-aptness measure seems *prima facie* to pertain to Independence. Yet, this becomes less apparent on a closer examination. For the truth-aptness task to measure Independence, for instance, participants should interpret each of the answer options as I described them above, namely “True” and “False” as *objectively* true or false and “opinion or attitude” as representing *subjective* opinions or attitudes.

However, these answer options are ambiguous. For example, the fact that a moral statement is treated as “True” or “False” does not mean that people believe that the statement is objectively true or false. It is also possible that they believe that it is subjectively true or false. Moreover, participants who select the option “False” may do so because they believe there are no true moral statements (see also Pölzler, 2018). Finally, the fact that participants select the option “opinion or attitude” does not necessarily provide evidence against FMO. Indeed, both opinions and the content of attitudes can be objectively true or false. That is, someone may believe that the statement ‘abortion is morally wrong’ is an opinion and yet objectively true (or false). The reason, for instance, why this person nevertheless does not select “True” or “False” is because s/he is careful about expressing his or her moral views to others. Consequently, research that used this particular formulation of the truth-aptness measure may have inadvertently introduced ambiguity (Goodwin & Darley, 2008; Wright, Grandjean & McWhite 2013; Wright, McWhite & Grandjean, 2014).

It is also unclear whether or not the disagreement task directly pertains to Exclusion. The reason is that the use of the term “surely” in the first answer option and the phrase “It could be that...” in the third answer option

may introduce a confounder variable (see also Pölzler, 2018). It may induce participants to interpret this question as being about how certain they feel about their responses. Moreover, in the disagreement task used by Goodwin and Darley (2008), the backgrounds of the parties involved in the moral disagreement are roughly presented as being similar to that of the participants.⁵ However, this means that we cannot assess whether or not participants judge in accordance with Exclusion because they are moral objectivists or for some other reason. Indeed, a cultural relativist can respond that at most one party is correct because they assume that the parties involved share their own cultural background (Sarkissian et al. 2012). Moreover, Pölzler (2018) notices that this measure also fails to account of the possibility that people might be non-cognitivists, which means that they believe that moral judgments are expressions of emotions or sentiments. Non-cognitivists do not think that it is possible to be correct or to be mistaken in a moral disagreement⁶. Consequently, there is no answer option that they can select. In light of these considerations, the disagreement task as used by Goodwin and Darley (2008) does not strictly measure Exclusion and provides no evidence for or against FMO.

2.3 A different experimental design to obtain high construct validity

In what follows, I will discuss research by Pölzler and Wright. They take significant methodological steps to ensure that research on folk metaethics is high in construct validity and avoids the problems inherent in measurement instruments used in previous research. However, I will argue that on closer inspection, different problems emerge.

As I interpret Pölzler and Wright, two assumptions underlie their basic approach. First, in order to avoid misunderstanding or different interpretations of metaethical concepts, Pölzler and Wright make the thesis of FMO as explicit as possible. Secondly, Pölzler and Wright make the answer options,

⁵ Participants are informed that another participant of the study disagrees with them about the content of a moral statement. It is likely that participants therefore assume that this participant is roughly similar to them.

⁶ With the exception of quasi-realists who do allow this (Blackburn, 1984, 1993).

and the different types of rejections of FMO, as explicit as possible. With this aim in mind, Pölzler and Wright decided to partly replicate previous research, namely by administering both abstract and concrete truth-aptness and disagreement tasks but in such a way that issues of construct validity are avoided. Moreover, they develop three additional measurement instruments, namely an abstract theory task, metaphor task, and comparison task. This was done to capture people's metaethical intuitions more comprehensively than previous research has done.

For the abstract truth-aptness task, participants read a description of what makes sentences truth-apt and not truth-apt. They were subsequently asked whether truth-aptness applies to moral statements or not. 73% of participants indicated that statements are truth-apt. For the concrete truth-aptness task, participants judged different moral statements and they were then asked to judge for each of those statements if those statements are truth-apt or not. On average, in 76% of the cases participants judged those sentences to be truth-apt. Consequently, people are largely inclined to believe that moral statements are truth-apt. Given that truth-aptness does not directly pertain to Independence, this does not provide evidence in favor of FMO. However, truth-aptness is a precondition for moral objectivism and the results here at least suggest that people, who have been elaborately informed about what truth-aptness entails, respond that they believe that moral sentences are true or false.

For the abstract disagreement measure, participants read a description of a disagreement between members of the same culture. They were then presented with different interpretations of that disagreement each of which corresponds to different metaethical views. One interpretation, for instance, is that one of the parties is right and the other person is wrong, which is a description of what Exclusion entails. Participants who provided a response that one of the parties is correct and the other person is wrong were subsequently presented with a moral disagreement between parties from different cultures (also see Sarkissian et al. 2012). This is to avoid the problem of misinterpreting

responses as favoring moral objectivism while people may actually be moral subjectivists. That is, people who are cultural relativists, and reject FMO, may appear to respond in accordance with Exclusion, and hence seem to support FMO, in cases where the parties to the disagreement belong to the same culture as they do.

Strikingly, 92% of participants deny that at most one party can be correct. Hence, the vast majority of people reject Exclusion. Given that this is a direct measure of Exclusion and one that excludes the possibility that people are moral subjectivists, this result appears to provide direct evidence against FMO.

For the concrete disagreement measure, participants were asked to consider each statement that they classified as moral and to then evaluate a disagreement between two people from the same culture. Those who responded in accordance with Exclusion then responded to a follow-up question. For this task, we similarly observe that on average 92% of participants respond in a way that fits with Exclusion and this therefore also appears to provide evidence against FMO.

What about the results on three additional abstract tasks that the authors administered? Do they similarly provide evidence against FMO? For an abstract theory task, participants read a description of the three main questions that determine one's position in the moral realism/anti-realism debate, namely "Do moral sentences intend to state moral facts?"; "If yes, do these facts exist?"; "And if yes, are they independent from what anybody thinks about them?" For the answer options, participants were provided with a range of descriptions of different metaethical theories such as realism, divine command theory, cultural relativism, individual subjectivism, error theory, and non-cognitivism. The rationale behind using these tasks, as described above, seems to be based on the assumption that the construct validity of research is improved if participants respond to questions that contain as maximally accurate descriptions of moral objectivism.

For example, the description of moral realism was as follows,

When a person says that something is morally right or wrong, good or bad, etc. she intends to state a fact. Such facts exist – and they are independent from what anybody thinks about them. For example, an action that is morally wrong is wrong no matter what anyone thinks. So it would still be wrong even if you yourself, or the majority of the members of your culture, thought that it is not morally wrong. (Pölzler & Wright, unpublished)

The results show, strikingly, that only 20% of participants select this option. The vast majority, 80% of participants, select an option that conflicts with Independence. Consequently, the results on the theory task also appear to provide evidence against FMO.

What about the results on a metaphor task? Participants were asked to select from a variety of metaphors describing different metaethical views. They first read the following description:

This task is about moral facts. Moral facts are facts about what is morally right or wrong, good or bad, virtuous or vicious, and so on. For example, it could be a moral fact that it is (or is not) wrong to break promises, or that the US has (or does not have) a duty to reduce their greenhouse gas emissions. Below moral facts are explained in terms of several metaphors. Which of these metaphors seem most appropriate to you. (Pölzler & Wright, unpublished)

The description for moral objectivism is:

“Moral facts are “discovered”. They can be discovered in the same way in which we discover other facts about the objective world”. (Pölzler & Wright, unpublished)

By using this description, this task only measures Independence insofar participants believe that moral facts are like scientific facts. That is, someone who believes that moral sentences can be true or false in an objective way but does not believe that they are a matter of empirical discovery may not select this option. If so, it is not clear whether this task measures FMO. For the metaphor task, only 13% of participants selected the metaphor that fits with moral realism. This seems to be a strong rejection of moral objectivism but it depends on whether people believe moral facts are like scientific facts.

Finally, participants responded to an abstract comparison task. In this task, participants read different descriptions of morality in comparison to other domains. With regard to morality, the comparison was made with science or mathematics and again a reference was made to discovering moral facts. As with the previous measure, negative answers do not have to entail that FMO is false. This depends on whether people believe that moral truth can be discovered like scientific truth. The results show that 10% of participants select the answer option that corresponds to moral realism.

In short, at the start of this section I remarked that the next logical step for research into folk metaethics was to construct studies with high construct validity. Pölzler and Wright aimed at having measurement instruments with high construct validity. Their approach was to avoid problems inherent to previous research and to be theoretically explicit about FMO and the different metaethical views that reject FMO.

What does this elaborate investigation tell us about FMO? First, the results suggest that most people believe that moral sentences are truth-apt. This does not mean that people support Independence, because they can believe moral sentences are subjectively true. Second, people overwhelmingly reject Exclusion both in the concrete and in the abstract disagreement task. This appears to be strong evidence against FMO. Third, when asked to evaluate different descriptions of morality (i.e. the Theory task), only 20% of participants select an option that corresponds to FMO. This also appears to provide strong evidence against FMO.

Given the responses on the Theory task and the fact that people appear to reject Exclusion, it appears that FMO is false. However, this conclusion is once again premature. FMO is only false if we take the evidence to bear directly on the commitments people have about the pertinent issues. In the next section, I will argue that this need not be the case.

3. Implicit and Explicit Moral Objectivism

In this section I will argue that people can have two distinct commitments toward moral truth and that FMO may apply to only one of them. This distinction is based on two structurally different types of psychological commitments that people can have, namely the distinction between implicit and explicit commitments. Philosophers have at this point not made a sharp distinction between those commitments but they have gestured at an intuitive difference between both. Moreover, there are some philosophers who seem to believe that only implicit commitments have evidentiary status with regard to FMO.

Brink (1989) claims, for example, that it is the presuppositions that underlie ordinary moral discourse and practices that show that people are moral objectivists. People's theoretical beliefs, on the other hand, have no bearing on the question of whether people are moral objectivists (Brink, 1989, p. 25). Similarly, Enoch writes that:

[W]hat is relevant is not the explicit metanormative beliefs – much less the explicit metanormative statements – of participants in normative discourse. What is relevant, rather, are the deep metanormative commitments embedded (perhaps implicitly) in normative discourse and practice themselves. The fact that many sophomores (and not only them) express some subjectivist or relativist metanormative intuitions thus has very little weight in assessing the commitments of normative discourse. (2005, p. 773, footnote 31)

Hence, if we follow Enoch's reasoning, findings of psychological studies are only relevant to the extent that they probe people's deep meta-normative commitments.

In a later paper, Enoch (2014) makes a similar distinction between implicit and explicit beliefs or commitments. He suggests that it is possible to be explicitly committed to a thesis while being implicitly (and unwittingly) committed to a different thesis instead. Enoch boldly claims: "You may think that you're a moral relativist or subjectivist – many people today seem to. But I don't think you are" (Enoch, 2014, p. 193). People may explicitly believe that they are moral relativists or moral subjectivists while they are, implicitly, committed to moral objectivism instead.

Björnsson (2012) draws a similar distinction between *commenting on* moral thought and discourse versus *being engaged in* moral thinking and debate. Indeed, Björnsson claims that "the primary task of metaethical theories is to account for this *engaged* behavior, rather than for what is in effect lay people's theoretical interpretations of it" (Björnsson, 2012, p. 9). I assume that what Björnsson describes as "engaged behavior" corresponds to Enoch's reference to deep meta-normative commitments embedded in normative discourse and practices. Just like Brink and Enoch, Björnsson believes that people's theoretical ideas about the status of morality are irrelevant.

The crucial question that we have to answer is then whether or not the research that I discussed in this paper measures implicit commitments or explicit commitments. This is a difficult question to answer because we first need a tool that we can use to identify implicit and explicit commitments. This is itself an empirical question that has not been properly investigated (but see Chapter 2). On my view, measurement instruments that directly and unambiguously explicate metaethical claims or views are most likely to elicit explicit commitments. That type of measurement instrument may cause people to theoretically speculate about the status of morality and provide a cognitive response instead of providing an answer that corresponds to how they ordinarily use moral language, how they think about moral issues, or what

they intuitively believe is the case. Therefore, implicit commitments should be measured indirectly.

In my view, studies that used scenarios or questions that refer to specific metaethical views or that used theoretically laden concepts or metaethical jargon may have activated an explicit mode of thinking in people (Nichols 2004; Sarkissian et al. 2012; Pölzler & Wright, unpublished). If so, people may have speculated about their answer instead of providing an answer that corresponds to what they implicitly believe is the case. Similarly, any research that used an experimental design that encourages people to explicitly reason about their answer may have measured explicit instead of implicit commitments. For example, participants in most research using truth-aptness and disagreement measures had to respond to a variety of moral statements. It is likely that the repetitive character inclines them to provide responses that do not reflect how they ordinarily think about these cases. More specifically, the repetitive nature of the experiment may have influenced people to provide reasoned instead of intuitive answers. Moreover, the sheer fact that people ordinarily do not explicitly reflect on whether or not a moral statement is true or false, or whether a moral disagreement concerns exclusionary content, may already have activated them to provide responses that reflect explicit commitments.

Consequently, it is hard to rule out that these studies measure theoretical conceptions of an explicit sort that metaethicists claim provide no insight on FMO. This means, of course, that we cannot straightforwardly interpret any findings until we know whether people's implicit commitments differ from the explicit commitments discussed above.

In my view, it is very likely that Pölzler and Wright's research design caused participants to mostly provide explicit commitments. This is because participants were provided with explicit theoretical conceptions of different metaethical concepts and views. Moreover, the distinction between normative ethics and metaethics was explained to participants, which may activate an explicit mode of thinking.

What seems reasonable to conclude, then, is that if FMO is a thesis about explicit commitments, it is false. But if FMO is a thesis about implicit commitments, we need to design empirical studies that measure these. At this point, we will therefore have to conclude that existing empirical research does not debunk FMO.

Conclusion

In this paper, I investigated whether or not there is empirical evidence for the philosophical assumption that people believe that moral judgments are objectively true or false, which I termed the thesis of folk moral objectivism (FMO). My analysis shows that many of the measurement instruments used in empirical research do not directly pertain to FMO. Moreover, those measurement instruments that appear to measure FMO may elicit explicit commitments while many philosophers believe that FMO concerns implicit commitments. To the extent that FMO pertains to explicit commitments, it is likely that FMO is false. To the extent, however, that FMO pertains to implicit commitments, it remains untested.

2.

FOLK MORAL OBJECTIVISM AND ITS MEASUREMENT⁷⁸

⁷ Most of this research was conducted during a stay as Visiting Assistant in Research at Yale University. I am deeply grateful to Joshua Knobe for his valuable support and advice during the development of this research and for the input of participants of his lab group meetings both at the Department of Philosophy and the Department of Psychology at Yale. In addition, I am grateful to Frank Hindriks, Jan Verplaatse, Bart Streumer, Daan Evers, Jacob Dijkstra, Mark Huismán, Kai Epstude, Gunnar Björnsson, Kari Wheeler-Reed, Kainika Quintelner, and Robbie Sutton and anonymous reviewers for their helpful comments and support. Finally, I would like to acknowledge the funding provided by the Research Foundation Flanders (FWO)—project nr. G.0683.13N.

⁸ This chapter is accepted as a paper at *Journal of Experimental Social Psychology*.

Abstract

Experimental philosophers and psychologists investigate whether people perceive moral judgments to be objectively true or false. Existing research focuses on a single dimension of ‘perceived objectivity’. The present research examines whether multiple dimensions of folk moral objectivity underlie moral judgments. It also examines whether such dimensions relate to perceived objectivity, tolerance, and people’s behavioral intentions to punish norm-violators. Exploratory factor analysis on twenty ethical items revealed three different ways of perceiving moral truth (Independent Truth, Universal Truth, Divine Truth), which each form reliable subscales (Study 1). This three-factor structure was supported by confirmatory factor analysis (Study 2). Each of the dimensions is differently related to perceived objectivity (Study 3). With respect to tolerance, perceived objectivity is a mediator in the relationship between perceiving moral truth as absolute or universal and tolerance (Study 4). With respect to a willingness to harm measure, Independent Truth is negatively related and Universal Truth is positively related, to people’s punitive attitudes toward norm-violators.

Keywords: Morality, Moral judgment, Moral objectivity, Tolerance, Willingness to harm Folk Moral Objectivism and Its Measurement

Introduction

In the past decade, psychologists and philosophers have started to investigate whether people perceive moral judgments to be objectively true or false by probing intuitions about moral objectivity. Existing research focuses on what is known as ‘perceived objectivity’⁹. This is often probed by two different questions, namely a truth-aptness task and a ldsagreement task (Goodwin and Darley 2008, 2010; 2012; Wright, Cullum & Schwab 2008; Wright, McWhite & Grandjean 2014). The former measures whether or not people believe that moral judgments are true or false. The latter measures how people respond to moral disagreements, namely whether or not one party is mistaken or that both can be correct. This emerging research literature has found large differences in objectivity ascriptions between individuals and between different moral issues, which has been termed metaethical pluralism (Wright, Grandjean & McWhite 2013).

Existing experimental research measures folk moral objectivity on a single dimension of perceived objectivity¹⁰. There are, however, good reasons to regard folk moral objectivity as multidimensional. First, people who perceive moral judgments as objective can have diverging reasons for doing so. Some people believe, for instance, that objective morality is constructed by the commands of a divine entity (Piazza & Landy 2013; Sarkissian & Phelan 2019; Yilmaz & Bahçekapili 2015). Others might regard moral judgments as true because they are derived from more basic moral truths (Kant 1785/1959). Second, someone who does not regard morality as objective might regard moral judgments as true relative to a culture (Harman 2012; Wong 2006). Alternatively, however, she might hold that moral judgments do not purport to describe an external states of affairs and are therefore neither true nor false;

⁹ A measurement scale provided by Forsyth (1980), his Ethics Position Questionnaire, does measure distinct ethical views. In light of Goodwin and Darley’s (2010) concerns of the relevance of the EPQ for measuring moral objectivity, this research attempts to develop a scale that can measure perceived objectivity judgments.

for example, as mere expressions of the subjective emotional states (Ayer 1936; Blackburn 1993)¹¹.

To contribute to existing experimental research on perceived objectivity, I present a scale for measuring folk moral objectivism (FMO), henceforth ‘the FMO-scale’. This scale accommodates universalism, absolutism, divine command theory, relativism, and what I call ‘no-truth’. The main innovation is that the FMO-scale allows for the possibility that folk moral objectivity has several dimensions. The scale is designed to test whether folk intuitions about moral objectivity are best captured in terms of a single psychological construct or by different psychological constructs. If folk attitudes toward moral truth and falsity are indeed best captured by multiple dimensions this has methodological implications for a wide range of experimental research in social psychology, including the ways in which experimental studies on folk moral objectivity have to be conducted.

I will first introduce different philosophical views that can underlie moral judgment and I discuss how the multidimensionality of folk moral objectivity has implications for different lines of research in social psychology. In Study 1, I use exploratory factor analysis to test whether there are common factors underlying twenty items measuring five distinct ethical views. This yields a three-factor structure and reliable subscales. In Study 2, I test whether the three-factor structure is supported by confirmatory factor analysis. Studies 3-5 serve to investigate whether the scale dimensions relate to perceived objectivity, tolerance, and willingness to harm.

¹¹ In analytical philosophy, this view on the semantics of moral statements falls under the heading of non-cognitivism. According to non-cognitivist views (e.g. emotivism or expressivism), moral statements do not intend to describe matters of facts but are perceived as ways of expressing non-cognitive mental states such as feelings, emotions or desires.

2. Beyond Perceived Objectivity

2.1. *Objectivity as a multi-dimensional construct*

So far, experimental research has focused on perceived objectivity, which has been a very fruitful endeavor. It is an empirical question, however, whether a single dimension of perceived objectivity captures folk attitudes towards the truth or falsity of moral judgments. People’s intuitions about this matter might vary on multiple dimensions. In this paper, I also consider universalism, absolutism, divine command theory, and the view that there are no moral truths.

Universalism, absolutism and divine command theory are different theories about why moral judgments are true. According to universalism, moral judgments are true only if they are based on universally binding moral norms that apply to anyone and everywhere (Hare, 1954; Quintelier, De Smet, & Fessler, 2013). An example of moral universalism can be found, for example, in the Universal Declaration of Human Rights. Article 1 of the declaration states that “all human beings are born free and equal in dignity and rights” and according to article 3 “everyone has the right to life, liberty and security of person” and so forth (The United Nations, 1948, my italics).

Moral absolutism goes beyond universalism in that it also holds that true moral judgments are derived from more basic moral truths. The underlying idea is that the core of morality is determined by a set of general rules and principles which all hold true, without exception (Wong 2006). Kant argued, for example, that moral obligations derive from the Categorical Imperative, which denotes the absolute moral requirement that one should “act only according to that maxim by which you can at the same time will that it should become a universal law” (1785/1959, p. 421). Kant famously argued that this implies that lying is prohibited even if you could save someone’s life by lying about her whereabouts.

Divine command theory is the view that whether an action is morally right or wrong depends on the commands of a divine being (Murphy, 1998; Quinn, 1978). In other words, true moral judgments are based on divine commands.

Those who support this theory regard religious texts and/or authorities as sources of moral knowledge. Morality, then, is what a divine being prescribes and acting morally consists of obeying divine commands. Each of these three positions entails perceived objectivity (see below for discussion). However, someone who denies divine command theory might still subscribe to absolutism. And someone who denies absolutism might still subscribe to universalism.

Just as there are different ways in which someone might affirm the objectivity of moral judgments, there are different ways in which someone might reject this. First, someone who denies moral objectivity might be a relativist and hold that the truth or falsity of a moral judgment is relative to cultures (Dreier, 1990; Harman, 1975; Wong, 2006). Second, people might also deny that there are moral truths. Perhaps they regard moral judgments as neither true nor false because they are expressions of emotions, which have no truth-value (expressivism; Ayer, 1936; Stevenson 1944, 1963). They could also believe that all moral beliefs are false (the error theory; Joyce, 2006; Mackie, 1977). I refer to this position as ‘the no-truth view of moral judgments’.

Some of the views discussed are closely related. As formulated above, absolutism entails universalism. This makes it rather unlikely that these positions reflect different dimensions of objectivity. One can, however, endorse universalism (and believe that norms have universal application) without subscribing to absolutism (and believing that such norms are derived from more basic moral principles). Furthermore, it is informative to see to what extent people’s responses reflect the degree to which positions are similar conceptually. It is possible that people’s endorsement of these views is predicted by a similar psychological mechanism.

All views either entail perceived objectivity or its denial. The point of the studies presented below is to determine whether there are important differences within the two camps. In order to tease this out, the FMO-scale does not include perceived objectivity as such. Instead, the relations between the above views and perceived objectivity are used to validate the scale. Note that the correlations are unlikely to be perfect. It is an open question whether the

different relationships that exist between these different views also constitute a psychological reality. The question is whether and how views about moral objectivity map onto human psychology. A related question is whether and how they relate to people’s tolerance toward morally divergent others and their willingness to harm others.

2.2 *The predictive power of folk moral objectivism*

Although people’s attitudes towards moral truth and falsity are interesting as such, it is also worth inquiring into what explains them and whether they make a difference to people’s tolerance judgments and behavioral intentions. As it turns out, there is large variation in perceived objectivity, both intra-personal and interpersonal, which has been termed meta-ethical pluralism (Wright, Grandjean & McWhite 2013). Moreover, perceived objectivity is related to social distance (Sarkissian, Park, Tien, Wright & Knobe 2011), religious background (Goodwin & Darley 2008; Sarkissian & Phelan 2019; Yilmaz & Bahçekapılı 2015), and age (Beebe, Qiaoan, Wysocki & Endara 2015; Beebe & Sackris, 2016). People with high scores on the personality trait of being open to experience tend to be moral relativists (Feltz and Cokely 2008). Those who have a competitive orientation towards argumentation are more often objectivists than those who have a cooperative attitude (Fisher et al. 2017). And Goodwin and Darley (2010) show that relativists score higher on disjunctive thinking.

As Sarkissian and Phelan (2019) observe, philosophers have suggested a relationship between religion and moral objectivity for more than two thousand years. Sarkissian and Phelan’s research shows that there is also an intricate psychological relationship between religion and moral objectivity. For example, one study shows that belief in a punishing God predicts people’s rejection of moral relativism. In a different study, the authors show that priming religious believers belonging to Abrahamic faiths with divinity concepts increases their objectivity ascriptions. Moreover, the researchers show that when people are generally prompted to believe in objective morality, they

are also more inclined to believe in a punishing God. Yilmaz and Bahçekaplı (2015) observed a similar relationship between religion and people's attitudes toward moral truth and falsity. They found that if people are primed with religious terms, they become increasingly objectivistic about morality, and when they are being primed with moral subjectivism they become less convinced of the existence of God. Consequently, religion and moral objectivity seem to be intertwined in particular circumstances, and it therefore makes sense to examine whether divine command ethics is a separate dimension of folk moral objectivity.

Existing research also suggests that folk attitudes towards moral truth make a difference regarding people's tolerance judgments (Goodwin & Darley, 2008; Nichols, 2004; Sarkissian et al. 2011; Wright, McWhite, & Grandjean, 2014). Perceived objectivity is associated with how comfortable people feel with interacting with morally divergent others (Goodwin and Darley 2012, Wright et al., 2014). Priming people with moral objectivism makes them twice as likely to donate to charity (Young and Durwin 2013), and priming them with moral relativism makes them more likely to cheat on an incentivized raffle and to engage in petty theft (Rai & Holyoak 2013; see also Tian 2008 about the relationship between moral relativism and moral behavior). This suggests that there is individual variation in how people construe moral judgments. And this opens up the intriguing possibility that those who agree about a particular moral issue (e.g. they both believe that abortion is morally wrong) may have different tolerance judgments toward others because they disagree about the status of moral judgments. In short, there is a surge of research on the effects of folk attitudes toward moral truth on a range of different variables. However, the relevance of research on folk moral objectivity goes beyond research on perceived objectivity.

2.3 *The psychological distinctiveness of moral attitudes*

The fact that people have different views on the status of moral judgments is something that distinguishes moral attitudes from non-moral attitudes.

The idea that people's treatment of moral issues is psychologically distinctive from their treatment of non-moral issues (i.e. conventions, preferences, taste, etc.) has been well established by existing research. From a young age, children have the basic capacity to distinguish moral violations (e.g. hitting someone) from conventional violations (e.g. talking out of turn) (Turiel, 1983, 1998; Smetana, 1981, 1983; Smetana and Braeges, 1990). People perceive violations of moral rules as less permissible and more serious than violations of conventional rules. Moreover, moral rules are perceived as authority-independent while conventional rules are perceived as authority-dependent (i.e., issued by decree of an authority figure or institution: moral rules cannot be changed in this way). Violations of moral rules are also perceived as generalizably wrong (i.e. wrong in other countries too) while the wrongness of conventional violations is perceived as local (i.e., wrong in a specific social situation, or culturally specific). Finally, justifications for moral rules are often given in terms of harm and welfare while justifications of conventional rules are given in terms of social acceptability.

The psychological distinction between moral and non-moral attitudes is also shown by the relationship that moral attitudes have with interpersonal tolerance. Haidt, Rosenberg, and Hom (2003) show that people are least supportive of moral diversity compared to other kinds of diversity. Moreover, perceiving an issue as moral instead of conventional increases people's intolerance toward morally divergent others, especially so in romantic or work contexts (Wright et al. 2008). Research by Skitka and colleagues shows that attitudes held with strong moral conviction ('moral mandates') increase people's preferred social and physical distance toward morally divergent others, it decrease people's goodwill and cooperativeness to resolve moral conflicts, and make people less willing to agree to procedural solutions to resolve disagreements (Mullen & Skitka 2006; Skitka et al. 2005; Skitka & Mullen 2002). What is crucial is that moral mandates explain unique variance beyond otherwise strong non-moral attitudes (Skitka et al. 2005). Hence, there is something special about moral attitudes, compared to non-moral attitudes, but what is it?

One thing that is special about moral attitudes according to research on the distinction between moral and conventional rules is that people generalize moral rules and violations. People are inclined to generalize moral rules to other contexts and situations, including different countries and cultures. This strongly resembles the philosophical view of moral universalism that I discussed above. Haidt, Rosenberg and Horn (2003, p. 6-7) explicitly build on the idea that people perceive moral judgments to apply universally.

What is special about moral attitudes compared to non-moral attitudes, according to Skitka et al. (2005) is that people perceive moral judgments as having universal application, that moral convictions refer to absolute beliefs that something is right or wrong, and that moral convictions are perceived as facts about the world. Citing Shweder (2002), Skitka et al. (2005) write that “[G]ood and bad are experienced as objective characteristics of phenomena and not just as verbal labels that people attach to feelings” (Skitka et al. 2005, p. 896-897).

The authors of the above studies unite different philosophical views, namely moral objectivity, universalism, and absolutism to explain the effects that moral attitudes have beyond non-moral attitudes. Given that these are distinct philosophical views, it is unclear whether they are a psychological conjoint, as assumed by existing research, or that they play different psychological roles. If folk moral objectivity is indeed a multidimensional matter, it becomes possible that these views play different psychological roles, and that they have different relationships to interpersonal tolerance and other variables of interest. It is therefore imperative that we investigate whether or not what is assumed to underlie moral judgments - namely attitudes toward their truth and falsity - fits on a single dimension of perceived objectivity or is best captured by multiple dimensions.

3. The Present Research

The goal of this paper is to investigate folk attitudes towards moral truth and falsity. Existing research measured perceived objectivity on a single dimension

and found large variance between individuals and between different moral statements. This metaethical pluralism in objectivity ascriptions may occur when individuals are presented with disagreement tasks about different moral issues. It is also possible, however, that folk attitudes toward moral truth are multi-dimensional and that individuals have different scores on each of the dimensions. This may, in turn, be differently related to perceived objectivity. To investigate this possibility, I set out to construct a multi-dimensional measurement scale of moral objectivity (Study 1). Subsequently, the three-factor structure that was found was tested in a confirmatory factor analysis (Study 2). The next three studies serve to establish the validity of the scale. Study 3 investigates the relationships between how people score on this scale on the one hand and perceived objectivity on the other. Study 4 concerns the relation to perceived objectivity and interpersonal tolerance. Study 5 concerns the question whether folk moral attitudes make a difference to people's behavioral intentions, namely their willingness to harm others.

4. Study 1: Development of the Measurement Scale

To investigate people's moral intuitions, the participants in this study were presented with a range of statements. The survey items are based on the positions discussed above: universalism, absolutism, divine command theory, relativism and no-truth. The point of the construction of a scale is to detect latent constructs. This requires exploratory factor analysis rather than principal component analysis (Tabachnik & Fidell, 2001). Participants' responses were analyzed using exploratory factor analysis with maximum likelihood extraction to test the underlying factor structure. It was then tested whether each of the three latent factors that were found form reliable scales that can be used as dependent and independent variables in subsequent studies. For this study and all studies that are part of this research and are presented below, all measures, manipulations and exclusions are reported. For all studies, sample sizes were determined before any data analysis.

4.1 Method

4.1.1 Participants

Four hundred fourteen participants were recruited via the online service Mechanical Turk and received \$0.50 for their time (188 female, Mage= 34, SD = 12.59). Amazon's Mechanical Turk provides an appropriate pool of research participants for research in psychology and the social sciences (Buhrmester, Kwang & Gosling, 2011; Paolacci & Chandler, 2014). Participants who did not complete the survey or failed to answer attention checks correctly (N = 10) were excluded from statistical analyses. The attention checks consisted of an item in the middle of the survey that instructed participants to remember the code word "Purple" and to rate "strongly agree" to that item. Participants were requested to fill out this code in a box on a new screen at the end of the study. Analyses were conducted on the remaining 404 participants.

4.1.2 Materials and Procedure

Participants received 20 items in a random order and were asked to rate the items on a six-point scale (1: Strongly disagree, 6: Strongly agree). The items consisted of statements about moral judgments that were developed on a variety of moral views, and were developed on the basis of philosophical literature on moral philosophy. They were further refined by consulting experienced philosophers and psychologists at three different universities and by brainstorm sessions at lab group meetings. A list of five different categories (universalism, absolutism, divine command theory, relativism, and no-truth) of items emerged (see Table 1). For example, items included "Without the existence

of God, nothing is truly morally right or wrong" (divine command theory) and "It is an illusion to think that anything is really morally true or false" (no-truth). This yielded twenty items in five different categories. Items were not reverse-scored for the purpose of factor analysis. Existing literature on scale construction warns against reverse scoring because it can lead to misinterpretation of items by participants and different types of measurement problems (Netemeyer, Bearden, & Sharma 2003; Swain, Weathers & Niedrich 2008). The results of the participants' responses were analyzed using exploratory factor analysis.

Table 1. Items and factor loadings Study 1 and scatterplot for the distribution of scores over the dimensions

Category	Item	Item label	Factor 1	Factor 2	Factor 3
No Truth	1	Other than what people believe, are brought up to believe, or want to believe about it, there are no facts about what is morally right and wrong	-.561	-.124	-.149
No Truth	2	All ideas about what is morally right and morally wrong are products of individuals, cultures, and communities and nothing more	-.751	.053	-.054
No Truth	3	What people believe to be morally right and wrong are merely social conventions that could have been different	-.718	.025	-.052
No Truth	4	It is an illusion to think that anything is really morally true or false	-.485	.044	-.293
Relativism	5	When two people have opposing beliefs about a moral issue, it is not necessarily the case that either or both are wrong	-.568	-.153	.029
Relativism	6	There is not one but many different answers to the question of what is morally right and wrong and these can be equally correct	-.724	-.077	-.025
Relativism	7	What is ultimately morally right and wrong is different for people with different moral views and from different cultures and societies	-.847	-.059	.094
Relativism	8	What is morally right and wrong is relative to the moral beliefs of an individual, culture, or society	-.861	.025	.081
Universalism	9	What is ultimately morally right or wrong is the same for all people at all times and places	.237	.201	.330

Universalism	10	Although people or cultures sometimes ignore moral concerns, moral norms apply anywhere and everywhere	-.071	.015	.750
Universalism	11	What is morally right and wrong for me here and now is also morally right and wrong for people elsewhere, even for people living in different countries and part of different cultures	.112	.128	.518
Universalism	12	Despite the diversity of moral views between individuals, cultures, and societies, there are moral norms that should apply universally	-.002	-.094	.767
Absolutism	13	Although people disagree about what is morally right and wrong, I believe in the existence of specific moral principles that can settle any moral disagreement	.110	.140	.621
Absolutism	14	Certain actions are morally wrong and they remain morally wrong even in the rare case that no one believes so	.013	-.017	.737
Absolutism	15	There are absolute moral rules that apply to all people, including those who do not acknowledge these principles	.050	.052	.687
Absolutism	16	There is, in all circumstances, one correct answer about what is the morally right thing to do	.105	.291	.478
DCT*	17	The correct answer to any moral issue can be found in a sacred book or text (for example, the Bible, the Qur'an, the Torah, or another)	.023	.782	.151
DCT*	18	The only actions that are ultimately morally right or wrong are those actions that God prescribes	-.047	.918	.068
DCT*	19	God is the only true source of knowledge about what is morally right or wrong	.037	.925	.000
DCT*	20	Without the existence of God, nothing is truly morally right or wrong	.007	.794	-.055
Eigenvalues			8.814	2.424	1.480
%Variance accounted for			44.07	12.12	7.40

Notes. *Divine Command Theory

4.1.3 Exploratory Factor Analysis

For all twenty items in the survey, exploratory factor analysis was performed with maximum likelihood extraction and using direct oblimin rotation, which made the assessment of latent constructs possible and also allowed factors to

be correlated (Fabrigar, Wegener, MacCallum, & Strahan, 1999). Sample size was determined by multiplying the number of items by factor 10, which suggests a sample of at least 400 participants (Velicer & Fava, 1998).

The data provides evidence for the existence of three separate factors. The significance of Bartlett's test for sphericity ($\chi^2_{5132.49} = (190)^2, p < 0.001$) and the value for the Kaiser-Meyer-Olkin test of .938 indicates excellent sampling adequacy (Kaiser, 1974; Tabachnik & Fidell, 2001). Eigenvalues, proportion variance explained, and factor loadings are reported in Table 1 and figure 1 visually depicts the distribution of scores between the different dimensions. As Table 1 shows, the items of these three factors loaded highly on only one factor and there are no cross-loadings. This indicates that the factors are distinct and it makes the three-factor structure interpretable and theoretically meaningful. Each factor contains items that strongly discriminate with items of a different factor. Eight items loaded on the first factor, four items loaded on the second factor, and eight items loaded on the third factor. The values of these indicators suggest that the measurement scale has excellent content validity. Consequently, the factors seem quite able to grasp the unobservable constructs under investigation. In light of the fact that item-development was based on theoretical considerations by consulting academic literature and experts in combination with the above results, it was decided to create a measurement scale out of these items and not to expand the list of items at this point.

Figure 1 Scatterplot for distribution of scores between the dimensions

Items expressing considerations related to relativism and no-truth compose the first factor. Because the items all loaded negatively on this factor, all of the items were reverse-coded. As a consequence, high scores on the scale indicate first, that there are moral truths and second, that they are independent of particular cultures. Because of this, I refer to this dimension as 'Independent Truth'. This scale had a high level of internal consistency (Cronbach's $\alpha = .90$). The second scale contains items that fall under the heading of divine command theory, which is the view that morality is based on a divine entity. This dimension is called "Divine Truth". The Cronbach's alpha score of .93 indicates high internal consistency. The third and final factor consists of items that express moral universalism and absolutism, which were combined into a scale labeled "Universal Truth". This scale has a similarly high level of internal consistency (Cronbach's $\alpha = .89$). Participants who score highly on this scale are more likely to support the idea that there are absolute moral norms that have universal application.

As allowed by the present analysis, and as is often the case for constructs in the social sciences, the factors correlate significantly. The Universal Truth and Independent Truth factors correlated moderately to strongly, $r(404) = .66$, $p < .001$, the Universal Truth and Divine Truth scores correlated moderately, $r(404) = .54$, $p < .001$, and the Universal Truth and Divine Truth scores also correlated moderately, $r(404) = .38$, $p < .001$.

4.2 Discussion

Study 1 presented participants with items about a range of moral views: universalism, absolutism, divine command theory, relativism and no-truth. The questions asked were whether people distinguish these positions and along which dimensions their intuitions about the status of moral judgments are structured. Perhaps unsurprisingly, divine command theory maps onto Divine Truth as a separate dimension. Universalism and absolutism turned out

to form one dimension, Universal Truth. Given how close they are conceptually, this is not very surprising either.¹²

The results also show that relativism and no-truth map onto one dimension. From a purely philosophical perspective this may seem a striking finding because relativism and no-truth seem logically inconsistent. That is, if there are no moral truths, then there are no relative moral truths either. However, from a psychological perspective this may be less surprising. This is because relativism and no-truth both reject the existence of a single objective truth. Hence, despite being philosophically distinct views, they consist of the same psychological construct in light of rejecting the idea of a single objective morality. The relativism and no-truth items fit on a single psychological dimension, Independence.

Another noteworthy result is that universalism and absolutism are not on the same psychological dimension as relativism and no-truth. In other words, Independent Truth and Universal Truth seem to be distinct dimensions. Moral judgments that are universally true will be true simpliciter as well as true independently of particular cultures. Even so, what is at stake in these two dimensions differs in that denying the former is different from denying the latter: people's beliefs about whether there are true moral judgments are correlated with, but independent from, their judgments about whether there are absolute moral principles or whether moral judgments are universally true. Yet existing research tacitly assumed that these views vary on a single psychological dimension (Bartels, Bauman, Cushman, Pizarro & McGraw 2016; Haidt, Koller & Dias 1993; Haidt, Rosenberg, and Hom; Skitka et al. 2005; Turiel, 1983, 1998). Hence, researchers were correct that views like objectivism, universalism, absolutism, and so forth, underlie morality. The present research contributes to this by showing that each of these views exist on different psychological dimensions.

¹² There were no cross-loadings between different factors, which suggest that the three factors are genuinely distinct. Of course, it is possible that some items for universalism and absolutism were too similar for people to be able to distinguish them. This is a possible limitation of this study and something to be improved in future research.

At first instance, the fact that Divine Command items load onto Divine Truth is not very surprising. However, theorists have often assumed that divine command ethics imply absolutism and universalism. Indeed, Sarkissian and Phelan acknowledge this by writing in their abstract that “[s]ome theorists contend that God is viewed as a divine guarantor of right and wrong, rendering morality universal and absolute”. The present research shows, however, that the Divine Truth dimension correlates with, but is distinct from, the Universal Truth dimension. This shows that, from a psychological perspective, adherence to divine command ethics does not imply a commitment to absolutism or universalism (though they are correlated).

In short, this study led to the development of a measurement scale, Folk Moral Objectivism (FMO). The results show that people’s intuitions about moral views vary on at least three distinct dimensions, namely Independent Truth, Divine Truth, and Universal Truth. The items of each of these dimensions form reliable subscales. These results show that moral views that are philosophically distinct do not have to be psychologically distinct. And they suggest that moral objectivity is best seen as a multi-dimensional construct. Of course, the specific selection of item-categories in this study may have constrained the possibility of discovering additional latent dimensions. Nevertheless, the statistical results on the present scale provide indications for a valid measurement scale. The question that I therefore address next is whether the FMO scale structure is confirmed by confirmatory factor analysis. If those results yield negative outcomes, I will consider additional philosophical views that may be part of folk moral objectivity.

5. Study 2: A confirmatory factor analysis of the three-factor structure

The next question is whether the factor structure, as revealed by exploratory factor analysis, will be confirmed by confirmatory factor analysis. Therefore, a separate data sample was collected to test whether this is the case. Confirmatory factor analysis is a structural equation modeling technique that allows one to test whether or not the shared variance of items can indeed be

explained by the three-factor structure of Independent Truth, Divine Truth, and Universal Truth.

5.1 Method

5.1.1 Participants

Four hundred ninety participants were recruited via the online service Mechanical Turk and received \$0.50 for their time (212 female; Mage = 34, SD = 11.80) and were presented with the same set of items as presented in Study 1. Participants who did not complete the survey or failed to answer attention checks (N = 9) correctly were excluded from statistical analyses. The attention checks consisted of an item in the middle of the survey that instructed participants to remember the code word “Rose” and to rate “strongly agree” to that item. Participants were requested to fill out this code in a box on a new screen at the end of the study.

5.1.2 Materials and Procedure

Participants received 20 items in a random order and were asked to rate the items on a six-point scale (1: Strongly disagree, 6: Strongly agree). The items consisted of statements about the status of morality that were developed on the basis of the five moral views discussed above (see Table 1) universalism, absolutism, divine command theory, relativism, and no-truth. For example, items included “Without the existence of God, nothing is truly morally right or wrong” (divine command theory) and “It is an illusion to think that anything is really morally true or false” (no-truth). The participants rated their agreement with each item. The results of the participants’ responses were analyzed using confirmatory factor analysis.

5.1.3 Confirmatory Factor Analysis

The confirmatory factor analysis was conducted in R 3.2.3, with the Lavaan package (Rosseeel, 2012) using maximum likelihood estimation to test the three-factor solution as found above. To examine the three-factor model as

revealed by exploratory factor analysis, CFA was used and the Comparative Fit Index (recommended: CFI > .90), the Standardized Root Mean Square Residual (recommended: SRMR < .08), and the Root Mean Square Error of Approximation (recommended: RMSEA < .08) were inspected (Bentler, 1990; Browne & Cudeck, 1993; Hu & Bentler, 1999). Results provided strong support for the three-factor model and indicated that it was a good fit (CFI = .928; SRMR = 0.053; RMSEA = 0.076). Taken together with the results of the exploratory factor analysis, we are more certain that each of the subscales of Independent Truth, Divine Truth, and Universal Truth reliably measures different latent dimensions containing items that form internally consistent subscales and that strongly discriminate between each other.

5.2 Discussion

The goal of Study 2 was to test whether the three-factor structure found by exploratory factor analysis in Study 1 would be confirmed by conducting a confirmatory factor analysis on a different data sample. The results of the CFA support the three-factor structure found by EFA. The scale thereby fulfills the basic requirements that we need for a proper scale: items were developed in extensive brainstorm sessions, all items load strongly on a single factor, there are no cross-loadings, subscales have high reliability, and CFA conducted on a separate dataset confirm the three-factor structure found by EFA. The question that I address next is how FMO relates to perceived objectivity.

6. Study 3: First validation study

In order to validate the FMO-scale, I now investigate whether it captures perceived objectivity. The scale was constructed to measure a wide range of intuitions concerning moral objectivity that are closely related to but distinct from perceived objectivity. Study 1 revealed three dimensions that underlie people's intuitions about moral objectivity. It might be that they predict people's perceived objectivity ascriptions as found in previously conducted experimental studies. If successful, this would confirm that the scale measures

what it purports to measure, and thereby support construct validity of the measurement scale.

To this end, I use the FMO-scale to replicate a previous study on perceived objectivity. Sarkissian et al. (2011) showed that people vary their ascriptions of objectivity when confronted with moral disagreements between people who have different moral frameworks. The question I ask is whether the dimensions of the measurement scale predict people's ascriptions of objectivity in a moral disagreement involving one party from the same culture and one party from a different culture (the other-culture condition of Sarkissian et al. 2011). In this study, this question was tested and thereby a first step toward validating the measurement scale was taken.

6.1 Method

6.1.1 Participants

Two hundred five participants were recruited via the online service Mechanical Turk and received \$0.50 for their time (114 female; Mage = 35, SD = 12.61). Participants who had previously participated in studies that involved the development of the measurement scale were excluded. Participants who did not complete the survey or failed to answer attention checks correctly (N = 2) were excluded from statistical analyses. The attention checks consisted of an item in the middle of the survey that instructed participants to remember the code word "Purple" and to rate "strongly agree" to that item. Participants were requested to fill out this code in a box on a new screen at the end of the study. Analyses were conducted on the remaining 203 participants. The sample size was based on an assessment of how the number of participants Sarkissian et al. (2011) used. A post-hoc power analysis using G*Power 3.1 (Paul, Erdfelder, Buchner & Lang, 2009) for a linear multiple regression design with a sample of 203 participants and an alpha of .05 indicated a statistical power of 99% to detect an effect size of 0.25, which is considered to be a small to medium effect size (Cohen, 1988).

6.1.2 Materials and Procedure

Each participant received the measurement scale and the other-culture condition from Sarkissian et al. (2011, Experiment 1) in counterbalanced order. In the other-culture condition, the participants were asked to imagine an isolated tribe of people called the Mamilions. They were told that the Mamilion tribe lives in the Amazon rainforests and has preserved a traditional warrior culture with different values from people in the surrounding society. Following this description, the participants received two questions to measure their responses to a moral disagreement concerning two different moral transgressions. Both transgressions involved canonical moral violations; one concerned the killing of a young child and the other involved the random stabbing of innocent passersby. The participants were then told that one of their neighbors thought that this act was morally wrong but that a member of Mamilion society thought that the act was morally permissible. For each transgression, the participants rated their agreement or disagreement with the sentence "Since your neighbor and the Mamilion have different judgments about this case, at least one of them must be wrong." The participants were asked to respond to this question on a scale of agreement from 1 to 7 (1: Disagree, 7: Agree)

6.2 Results

Each participant was given a score representing his or her mean judgment of the two moral transgressions. First, correlations between the dimensions of the measurement scale and the composite score of the other-culture judgments were calculated. The other-culture score positively correlated with Independent Truth, $r(203) = .51$, $p < .001$, Universal Truth, $r(203) = .54$, $p < .001$, and Divine Truth $r(203) = .17$, $p = .02$. Subsequently, the other-culture score was regressed on the three dimensions and results show that the three predictors explained 34.2% of the variance ($R^2 = .342$, $F(3, 199) = 34.51$, $p < .001$). Independent Truth ($b = .45$, $t(3, 199) = 2.98$, $p = .01$), Universal Truth ($b = .79$, $t(3, 199) = 4.89$, $p < .001$), and Divine Truth ($b = -.20$, $t(3, 199) = -2.15$, $p = .03$) predicted vignette scores.

6.3 Discussion

This study reveals that the dimensions of the measurement scale are each associated with perceived objectivity. Independent Truth and Universal Truth are positively related to this construct. That is, those who score highly on Independent Truth or Universal Truth tend to say that at least one of those involved in a moral disagreement has to be wrong. Divine Truth, in contrast, is negatively related to perceived objectivity.

The findings of this study validate the FMO-scale in that it successfully predicts perceived objectivity as measured by the other-culture condition of Sarkissian et al. (2011). These results suggest that the Sarkissian probe did not simply measure a one-dimensional construct but instead tapped into three different dimensions.

Divine Truth is as such positively correlated to perceived objectivity. However, a multiple regression revealed a negative but small relationship between Divine Truth and the other-culture variable when controlling for Independent Truth and Universal Truth. For someone who adheres to divine command ethics, the results here suggest that the positive association between Divine Truth and perceived objectivity results from an endorsement of Independent Truth and/or Universal Truth. This relationship was not predicted. Indeed, Goodwin and Darley (2008) show a positive association between grounding one's morals in divine command ethics and perceived objectivity. Yilmaz & Bahçekapili show that there is an association between religious concepts and moral objectivity. Moreover, Sarkissian and Phelan (2019) show that followers of Abrahamic faiths are most likely to endorse moral objectivism and that it is specifically people's beliefs in God's punishing characteristics that predict moral objectivity. Those are interesting findings, but those studies did not distinguish between Independent Truth, Universal Truth and Divine Truth. It is therefore possible that the relationship that was found between religion and moral objectivity is merely correlational and disappears when controlling for Independent Truth and Universal Truth.

It might be that people have different views on whether or not the commands of a divine entity apply to other cultures. If that is the case, people who score high on Divine Truth may respond as if morality is relative because they believe that the divine commands issued by God apply to their own culture and not necessarily to members of different cultures. Indeed, God may even have different commands for members of different cultures. Alternatively, it is possible that people recognize that other cultures have different gods and that those gods may issue different commands. As a result, moral truth is relative to those different cultures.

A yet different possibility is that an individual differences variable explains the specific relationship found here. The study shows that people who are (more) committed to a divine command ethics tend to refrain from judging moral disagreements as if only one person is correct. Thus, although they believe that moral truths are based on divine commands (as the scale items measure), they refrain from explicitly judging that others must be mistaken. The reason for this may be that people who are committed to divine command ethics may also be the kind of people that refrain from judging what others should do or think. Indeed, they may believe in a very personal form of free will and moral responsibility - that is, it is ultimately God who will judge who was right and who was wrong. Among like-minded thinkers, they perhaps may judge that a certain moral truth exists, but when they are asked about this in a more detached forum, like these studies, they may refrain from making judgments. Alternatively, Saroglou (2011) proposes that there are four distinct dimensions of religion (believing, bonding, behaving, and belonging) that may express themselves differently in different cultures. It is possible that people's view about whether or not the commands of a divine entity apply to other cultures depends on their psychological reliance on each of the dimensions identified by Saroglou (2011). Relatedly, the scores of individuals on Divine Truth may then be culturally variable (see also Graham & Haidt, 2010; Graham, Meindl, Beall, Johnson & Zhang, 2016). Consequently, there are a variety of possibilities to explain the effect here and this would be an interesting question for future research.

7. Study 4: Second validation study

The aim of Study 3 was to further validate the measurement scale by replicating a previous study on perceived objectivity and interpersonal tolerance. Goodwin and Darley (2012) found that people who hold that at least one of two people who morally disagree must be mistaken tend to be less tolerant toward morally divergent others. In this study, I set out to replicate Goodwin and Darley's (2012) study. They operationalized tolerance in terms of how comfortable a participant would be to have a long-term guest who disagreed with them. Just as Study 2, this experiment can be used to check whether the dimensions of the scale predict people's objectivity ascriptions. Similarly, this study uses Goodwin and Darley's cases to test whether the three dimensions of the FMO-scale - to wit Independent Truth, Universal Truth, and Divine Truth - predict perceived objectivity.

The main goal of this study, however, is to investigate whether the dimensions of the scale predict people's tolerance toward morally divergent others. There is reason to expect each of these dimensions to be predictive of intolerance. Those who submit to Independent Truth, Universal Truth, or Divine Truth might assume that they know what is right and wrong and might be critical of those who have moral beliefs that they regard as mistaken. However, this need not be the case. Belief that moral truths exist can be combined with the belief that it is not always easy to know those truths. Even so, given the positive relation Goodwin and Darley found, I expect that at least Independent Truth and Universal Truth are predictive of intolerance. The fact that Divine Truth turned out to be negatively correlated to perceived objectivity in Study 3 suggests that it is an exception.

In this study I will also investigate the differences between perceived objectivity and the dimensions of the measurement scale with regard to tolerance. Goodwin and Darley's study have shown that perceived objectivity is associated with tolerance. The previous study (Study 3) shows that the dimensions of the measurement scale are associated with perceived objectivity. The question is then what associations exist between perceived objectivity, the dimensions of the measurement scale, and tolerance. I will test whether per-

ceived objectivity and/or the dimensions of the measurement scale predict tolerance when controlling for each other.

7.1. Methods

7.1.1 Participants

Three hundred fifty participants were recruited via the online service Mechanical Turk and received \$0.50 for their time (174 female; Mage = 35, SD = 12.17). Participants who previously participated in studies that involved the development of the measurement scale were excluded. Participants who did not complete the survey or failed to answer attention checks correctly (N = 11) were excluded from statistical analyses. The attention checks consisted of an item in the middle of the survey that instructed participants to remember the code word "Yellow" and to rate "strongly agree" to that item. Participants were requested to fill out this code in a box on a new screen at the end of the study. Analyses were conducted on the remaining 347 participants. A large sample size was chosen in order to have sufficient statistical power. A post-hoc power analysis using G*Power 3.1 (Faul, Erdfelder, Buchner & Lang, 2009) for a linear multiple regression design with a sample of 339 participants and an alpha of .05 indicated a statistical power of 100% to detect an effect size of 0.25, which is considered to be a small to medium effect size (Cohen, 1988).

7.1.2 Materials and Procedure

Each participant received the FMO-scale as well as measures of perceived objectivity and tolerance in a counterbalanced order. The measure of objectivity and tolerance was the same as that used by Goodwin and Darley (2012, Experiment 1). For the measure of objectivity and tolerance, each participant received six different scenarios involving moral issues, in a random order:

1. **Steal wallet.** Jason is saving up for an iPod, but he is getting impatient that it is taking so long to have enough money. After he has finished dinner at a local restaurant one evening, he notices that another customer has left their wallet

behind on the table next to him. He is able to look inside the wallet discreetly, and finds \$200 in cash. He takes the \$200, and leaves the restaurant. Rate the extent to which you agree with the claim that Jason's actions are morally wrong.

2. **Punch.** After a very difficult day at work, Frank goes to his local bar to watch his favorite team. As soon as Frank sits down, he overhears a fellow patron make disparaging comments about Frank's team to the bartender. Frank immediately walks over to the person who made the comment, and punches him off his bar stool. Rate the extent to which you agree with the claim that Frank's actions are morally wrong.

3. **False alibi.** One of Megan's best friends is being charged with murder. Megan is convinced that he is innocent, although she does not know what he was doing on the night of the alleged murder. Without having been asked, Megan provides a false alibi to the police for her friend, claiming that she was with him on the night of the night of the alleged murder. Rate the extent to which you agree with the claim that Megan's actions are morally wrong.

4. **Urinate on memorial.** Tom is out with his friends one night and has been drinking. As they are walking home, they encounter a memorial for victims of 11/9, with flowers laid at the base of it. Tom wants to impress his friends, and so he decides to vandalize the memorial. He urinates on the memorial and on the flowers. Rate the extent to which you agree with the claim that Tom's actions are morally wrong.

5. **Nazi salute.** Mike is a professional sportsman. He is playing in a match against a team that is known to have a large Jewish support-base, and these opposition supporters are heckling him. He responds by turning to these supporters, mimicking Adolf Hitler's mustache, and giving them a Nazi salute. Rate the extent to which you agree with the claim that Mike's actions are morally wrong.

6. *Burn flag.* Amy is a high school history teacher. She has become increasingly dissatisfied with her government's foreign policies, and wants to communicate that to her students. She decides to do this by burning a US flag in front of them. Rate the extent to which you agree with the claim that Amy's actions are morally wrong.

The participants read each scenario and were then asked to rate the extent to which they agreed that the person's actions were morally wrong on a six-point scale (1: strongly disagree, 6: strongly agree). The participants were then asked two objectivity questions and a tolerance question. For the first objectivity question, the participants were asked whether there was a correct answer to whether the moral claim was true (1: no correct answer, 6: definitely a correct answer). For the second objectivity question, the participants were asked how they would interpret a moral disagreement with regard to the moral claim (1: Neither of us needs to be mistaken, 6: The other person is clearly mistaken). Finally, for the tolerance question, the participants were asked how comfortable they would be to have a long-term guest who disagreed with them (1: Extremely uncomfortable, 6: Extremely comfortable).

7.2 Results

The two objectivity measures were combined to generate a composite measure of moral objectivity by adding up the scores and dividing it by the number of variables. Each participant was then given an objectivity score representing the mean of the judgment for the six moral transgressions. I also constructed a composite score of the tolerance variable for each of the six moral transgressions by adding up scores and dividing it by the number of variables. Subsequently, correlations between the dimensions of the measurement scale and the composite score of perceived objectivity were calculated. Independent Truth, $r(338) = .37$, $p < .001$, Universal Truth, $r(338) = .42$, $p < .001$, and Divine Truth, $r(338) = .17$, $p < .001$, each correlated significantly with the composite objectivity score. Correlations were also calculated for the tolerance measure. Independent Truth, $r(339) = .18$, $p = .001$, and Uni-

versal Truth, $r(339) = -.21$, $p < .001$, significantly correlated with tolerance, but Divine Truth did not, $r(339) = -.08$, $p = .163$.

The next step was to investigate the relationship between the dimensions of the measurement scale, perceived objectivity and tolerance. To this end, I computed several linear mixed effect models in R 3.2.3 using the lme4 package (Bates, Maechler, Bolker & Walker 2015). A mixed effects model analysis makes it possible to account for the cross-nestedness of observations in respondents and scenarios through the inclusion of random effects (e.g., Gelman and Hill 2007). In the first model, the relationships between Independent Truth, Universal Truth, and Divine Truth and the composite objectivity score were investigated. The results in Table 2 show that Independent Truth and Universal Truth are both positively related to the composite objectivity score. The model controls for age, gender, and nationality. Comparison to an otherwise identical model without random scenario effects (results not shown) reveals that the latter significantly improved model fit ($\chi^2 = 347.45$, $df = 1$, $p < .001$). Figure 2 depicts the estimated random scenario effects for the model in Table 2.

Table 2

Mixed Effects Model with perceived objectivity as dependent variable

Variables	Coefficients	S.E.	T-ratio
Fixed Effects:			
Intercept	5.92	.57	10.39
Independent Truth	.28	.13	2.15
Universal Truth	.65	.13	5
Divine Truth	-.09	.07	-1.29
Random Effects:			
Respondents	1.47		
Scenarios	1.04		
Residual	2.15		
Deviance	9537.3		

Notes. Controlled for Age, Gender, and Nationality

Figure 2 Plot of random effects of scenarios for composite objectivity score

Subsequently, I examined the relationship between the dimensions of the measurement scale and the tolerance variable. The results presented in Table 3 show that Universal Truth is significantly related to tolerance while Independent Truth and Divine Truth are not, again controlling for age, gender, and nationality. Again, comparison to an otherwise identical model without random scenario effects (results not shown) shows that these effects significantly improve model fit ($\chi^2 = 138.59$, $df = 1$, $p < .001$).

Table 3
Mixed Effects Model with tolerance as dependent variable

Variables	Coefficients	S.E.	T-ratio
Fixed Effects:			
Intercept	3.83	.28	13.68
Independent Truth	-.03	.08	-.375
Universal Truth	.23	.08	2.88
Divine Truth	-.09	.05	-1.8
Random Effects:			
Respondents	.95		
Scenarios	.36		
Residual	1.17		
Deviance	7100.7		

Notes. Controlled for Age, Gender, and Nationality

The next step was to include the fixed effect of the composite objectivity score in the model of Table 3. The results presented in Table 4 reveal that in this model only the composite objectivity score is significantly related to tolerance. In contrast to the model without perceived objectivity, the relationship with Universal Truth is not statistically significant. What these results indicate is that the composite objectivity score is a mediator in the relationship between Universal Truth and tolerance (see figure 4). As in the previous models, comparison to an otherwise identical model without random scenario effects shows the latter improve model fit significantly ($\chi^2 = 19.94$, $df = 1$, $p < .001$). Figure 3 shows the estimated random scenario effects of the model in Table 4.

Table 4 Mixed Effects Model including composite objectivity score with tolerance as dependent variable

Variables	Coefficients	S.E.	T-ratio
Fixed Effects:			
Intercept	5.21	.24	21.71
Independent Truth	.03	.08	-.375
Universal Truth	-.07	.08	-.875
Divine Truth	.02	.04	.5
Objectivity	-.23	.01	-.23
Random Effects:			
Respondents	.89		
Scenarios	.15		
Residual	1.06		
Deviance	6682.7		

Notes. Controlled for Age, Gender, and Nationality

Figure 3 Plot of random effects of scenarios for tolerance

Figure 4 Coefficients for the relationship between Universal Truth and Tolerance with Perceived Objectivity as mediator

7.3 Discussion

As in Study 3, Independent Truth, Universal Truth, and Divine Truth, are positively correlated to perceived objectivity. A mixed effects analysis shows that it is particularly Independent Truth and Universal Truth that predict perceived objectivity. Furthermore, a similar analysis shows that Universal Truth is negatively related to tolerance and that perceived objectivity is a mediator in this relationship. This study thereby reveals that the dimensions of the measurement scale explain distinct variance in perceived objectivity and tolerance. In particular, this study presents a further step toward validation of the FMO-scale by showing that the dimensions are not only related to perceived objectivity but also to people's tolerance of others who morally disagree with them.

Just as in Study 3, results indicate that the relationship between Divine Truth and perceived objectivity seems less straightforward as results of previous research suggested (Goodwin & Darley 2008; Sarkisian & Phelan 2019; Yilmaz & Bahçekapili 2015). Divine Truth and perceived objectivity correlate positively but the relationship disappears when Independent Truth and

Universal Truth are taken into account. Given that a similar result was obtained in Study 3, it seems that this relationship is genuine. Future research should investigate whether or not any of the explanations given in the discussion of Study 3 explain this result.

The results here also show that Universal Truth, when controlling for Independent Truth and Divine Truth, decreases people's tolerance of morally divergent others. More specifically, perceived objectivity is a mediator between Universal Truth and tolerance. People's commitments to Universal Truth are positively related to perceived objectivity while perceived objectivity is negatively related to tolerance. This suggests that perceived objectivity works as a mechanism between Universal Truth and tolerance.

Skitka et al. (2005) show that the strong moral convictions that people have decrease their tolerance of those who have different moral convictions. It increases people's social and physical distance from others, and it decreases their goodwill and cooperation to resolve conflicts. In their seminal paper, Skitka et al. (2005) explicitly assume that moral judgments are perceived as being objective, universal, and absolute. The present studies suggest that specifically absolutism and universalism (represented here by Universal Truth) are associated with tolerance. Hence, Skitka et al. (2005) correctly identified that different views underlie moral judgments and that these can potentially explain associations with tolerance. The present research contributes by showing that Universal Truth is particularly associated with tolerance. It is therefore possible that people's scores on Universal Truth explain the results found by Skitka and colleagues. If this is correct, future research should distinguish between the different views that underlie moral judgments and hence at least take into account Independent Truth and Universal Truth.

In short, this study shows that Independent Truth and Universal Truth explain distinct variance in perceived objectivity. Moreover, it shows that perceived objectivity is a mediator in the relationship between Universal Truth and tolerance. Previous research assumed that distinct views underlie moral judgments and the present study contributes by showing that these views are

differently related to tolerance. An important question now is whether similar relationships exist between the dimensions of the measurement scale and measures of behavioral intentions that are different than tolerance.

8. Study 5: Willingness to harm

Results of Study 4 suggest that Universal Truth is associated with people's tolerance toward those who morally disagree with them. In Study 4 tolerance was measured as the degree to which people are comfortable with having someone who morally disagrees with them as a long-term guest in their house. The aim of the present study is to determine whether the relation that is found between people's attitudes towards moral objectivity and intolerance also extends to variables beyond intolerance, namely people's willingness to harm others.

As discussed above, research by Skitka and colleagues (Skitka et al. 2005) show that attitudes held with strong moral conviction, which they term moral mandates, decreases people's tolerance, goodwill, and cooperativeness to those who morally disagree with them. In their research, they do not distinguish between the views that underlie moral judgments. The present research investigates whether Independent Truth, Universal Truth, and Divine Truth have different relationships with a behavioral intention measure that extends beyond tolerance.

The general hypothesis is therefore that Independent Truth, Universal Truth, and Divine Truth are differently associated with people's willingness to harm others. This study was preregistered at Aspredicted.org (#1531:<http://aspredicted.org/blind.php?x=6rv64r>)

The willingness to harm measure was chosen to test whether folk attitudes toward moral truth and falsity have relationships with variables beyond tolerance. The tolerance variable used in Study 4 measures the degree of physical and social distance people desire from people who have different moral convictions than they do. Tolerance is one specific response to people who

have different moral convictions. A different question pertains to how people respond to situations in which other people violate specific moral norms.

People who violate moral norms are ordinarily punished for their behavior. This punishment can take place with the aim of incapacitation or deterrence. People often experience moral outrage if they perceive others violating a moral norm. If they decide to punish the offender they may do this to incapacitate further moves, they may do this to deter people from violating those norms in the future, or they may do this to signal to the offender why he or she is being punished (Darley & Pittman 2003; Gollwitzer 2009). Regardless of motive, it may be the case that people's perception of moral truth relate to people's willingness to harm those who violate moral norms.

For instance, if people adhere to Independent Truth, Universal Truth, or Divine Truth, they may believe that they have knowledge of what is morally right and wrong. As a result, they adopt a critical stance toward those who violate moral norms irrespective of their scores on the dimensions of the measurement scale. Alternatively, each of the dimensions may be differently (i.e., positively and negatively) related to willingness to harm. Indeed, the perception that moral truth is universal and absolute, as represented by Universal Truth, might induce people to respond in a resolute manner to those who violate moral norms. Perceiving moral truth as absolute and universal might entail a justification for setting other people straight. Alternatively, believing in the possibility of a single objective moral truth (i.e. rejecting no-truth and moral relativism), without believing that moral truth is absolute or universal, might inhibit people to respond violently. This might be the case because those who violate moral norms are merely perceived as being led astray by mistaken beliefs about the single objective truth and not as violating absolute or universal moral norms. A specific hypothesis here is therefore that Universal Truth increases while Independent Truth decreases people's willingness to punish norm-violators.

In this study, participants received the measurement scale and seven different scenarios in counterbalanced order. Each of the scenarios is described

as involving a party violating a moral norm and a different party as having the option of responding harmfully. Participants were asked to indicate whether or not they would favor a harmful response. The scores on the measurement scale were used to predict whether people have a willingness to harm others but no specific direction was hypothesized.

8.1 Methods

8.1.1 Participants

Four hundred ninety-three participants were recruited via the online service Mechanical Turk and received \$0.50 for their time (257 female; Mage= 34). Participants who previously participated in studies that involved the development of the MMS were excluded. Participants who did not complete the survey or failed to answer attention checks correctly ($N = 12$) were excluded from statistical analyses. The attention checks consisted of an item in the middle of the survey that instructed participants to remember the code word "Orange" and to rate "strongly agree" to that item. Participants were requested to fill out this code in a box on a new screen at the end of the study. Analyses were conducted on the remaining 480 participants. I chose for a relatively high number of participants for this study to be certain that there is sufficient statistical power to detect a small to medium effect size. Post-hoc power analysis using G*Power 3.1 (Faul, Erdfelder, Buchner & Lang, 2009) for a linear multiple regression design with a sample of 480 participants and an alpha of .05 indicated a statistical power of 100% to detect an effect size of 0.25, which is considered to be a small to medium effect size (Cohen, 1988).

8.1.2 Materials and procedure

Each participant received the measurement scale and seven different scenarios in counterbalanced order. The scores on the measurement scale were used to predict whether people are willing to harm others, which was measured on a scale from 0 (do not use violence) to 100 (use violence). For example, one scenario concerns the President of the United States deciding about what to

do to stop a violent terrorist group. The President is described as considering using force and bombing the terrorist group to stop them. A different scenario describes a gay couple in a restaurant minding their own business and then suddenly being targeted by two bullies who verbally harass them about their sexual orientation. In the first scenario, the scale ranged from 0 ("Don't bomb them") to 100 ("Bomb them"). In the second scenario, the scale ranged from 0 ("Don't pull them off") to 100 ("Pull them off the barstool"). The full set of 7 scenarios:

1. **President.** The president of the United States is deciding about what to do to stop a terrorist group. The terrorist group has recently committed terrifying brutalities to innocent people and the President of the United States is horrified by this. He is considering to use force and bomb them.

2. **Gay couple.** A gay couple in a restaurant was minding their own business but then were suddenly targeted by two bullies who started calling them names and otherwise verbally harassing them about their sexual orientation. Two onlookers, at the other end of the bar, are appalled by the way that these bullies are treating the gay couple. They are thinking of teaching the bullies a lesson by pulling them violently off their bar stools.

3. **Dog.** Jason is walking his dog and observes another dog owner hitting his dog with a belt. The dog is clearly in pain and whimpers loudly. Jason is thinking of making him understand how the dog feels by hitting the dog owner once with the same belt.

4. **College.** Jack, a senior in college, and Jessica, a freshman, are siblings and both in the same college. Jessica has told Jack that she was sexually harassed by one of her male classmates. Jack is considering visiting the classmate and setting him straight by putting him in an armlock and hurting him.

5. **Football.** Some people are playing football. A few people on the sideline start shouting racial slurs at the black players and they even throw bananas on the field. Most of the players on the field believe that this is the wrong way to treat people and are thinking about violently throwing them out of the stadium.

6. **Christians.** A group of Christians is holding a public prayer session in a park. They are unexpectedly interrupted by a person who ridicules the Christian faith by shouting loud and disrespectful remarks about the bible in order to offend them and interrupt their prayer. A jogger in the park who happens to run by is offended and thinks about pushing the person into the fishpond.

7. **Veterans.** At the end of Veterans Day, a group of veterans has come to a restaurant to have dinner together. While the veterans are having dinner, two other guests in the restaurant start mocking the veterans and make insulting remarks about the American army. To make things worse, they even use a t-shirt with the American flag to clean one of the dishes. The veterans are thinking of inviting these people to a fight outside to teach them a lesson.

8.2 Results

The scores on the willingness to harm measure for each of the seven scenarios were combined to generate a composite measure of willingness to harm. Therefore, each individual participant had a unique average score on willingness to harm.

Subsequently, correlations between the dimensions of the FMO-scale and the composite score willingness of harm measure were calculated. Independent Truth, $r(480) = -.118$, $p < .001$, and Universal Truth, $r(480) = .102$, $p < .001$, dimensions each correlated significantly with the willingness to harm score but Divine Truth, $r(480) = .02$, $p = .653$, did not.

The next step was to investigate the relationship between the dimensions of the measurement scale and willingness to harm. To investigate these relationships,

I conducted a linear mixed effect analysis in R 3.2.3 using the lme4 package. The results presented in Table 5 show that controlling for age, gender, and nationality Independent Truth is negatively related to willingness to harm while Universal Truth is positively related to willingness to harm.

I compared the model from Table 5 to an otherwise identical model without random scenario effects, revealing significant effects of the latter ($\chi^2 = 504.69$, $df = 1$, $p < .001$). Figure 5 depicts the estimated random scenario effects in the model for willingness to harm.

Table 5

Mixed Effects Model with willingness to harm as dependent variable

Variables	Coefficients	S.E.	T-ratio
Fixed Effects:			
Intercept	42.19	6.19	6.82
Independent Truth	-5.6	1.24	-4.52
Universal Truth	5.87	1.41	4.16
Divine Truth	-.33	.86	.38
Random Effects:			
Respondents	Std. Dev.		
Scenarios	21.95		
	10.61		
Residual	23.48		
Deviance	31964.3		

Notes. Controlled for Age, Gender, and Nationality

Figure 5 Plot of random effects of scenarios for willingness to harm

8.3 Discussion

I hypothesized that Independent Truth and Universal Truth can be differently related to willingness to harm. That is, perceiving moral truth to be absolute and universal may induce people to respond resolutely to violations of moral norms. Additionally, believing that there is a single objective truth (rejecting moral relativism and no-truth), without believing moral absolutism or universalism, may inhibit people to respond violently. The results of this study show that Universal Truth and Independent Truth are indeed differently related to willingness to harm, respectively positively and negatively.

The results of this study are important for a variety of reasons. First, Independent Truth and Universal Truth explain distinct statistical variance in willingness to harm. This implies that Independent Truth and Universal Truth are psychologically distinct dimensions, which further validates the FMO scale. Second, the fact that Independent Truth and Universal Truth pull into different directions when it comes to willingness to harm suggests that these dimensions fulfil distinct psychological roles. Third, these results

provide a novel perspective on research on strong moral attitudes and their relationship with different types of intolerance (e.g. Skitka et al. 2005). It is possible that examining people's strong moral attitudes, while controlling for Independent Truth and Universal Truth, reveals different relationships with the aforementioned variables.

9. General Discussion

Five studies suggest that folk moral objectivity is a multi-dimensional phenomenon. The results of these studies provide new insights into how people think about the objectivity of morality, and they provide a novel tool for measuring people's intuitions about this, the FMO-scale.

Study 1 shows that moral objectivity can be measured on distinct dimensions. People's responses to twenty items from five different categories (representing the philosophical views of universalism, absolutism, divine command theory, relativism, and what I call 'no truth') revealed three distinct psychological constructs. The dimension of Independent Truth captures the view that moral judgments that are true independently of the group or culture to which those who form the judgments belong. People who endorse this dimension reject the idea that there are no moral truths or that moral truth is relative and hence share the idea that there is a single objective truth. The dimension of Universal Truth represents the view that there are absolute moral norms that are universally binding. Finally, the dimension of Divine Truth concerns the view that what is morally true or false depends on the existence of a divine entity and that moral knowledge is revealed in divine books and religious texts.

Studies 2 and 3 show that Independent Truth and Universal Truth, but not Divine Truth, are positively related to perceived objectivity. Studies 3 and 4 also reveal a negative relationship between Universal Truth and interpersonal tolerance. Study 4 shows that the dimensions of Universal Truth and Independent Truth pull people's willingness to harm others in different directions, which indicates that these dimensions play distinct psychological

roles. Whereas high scores on Universal Truth are associated with a relatively high willingness to harm others, those who score highly on Independent Truth express less willingness to harm others.

9.1 Relationship to previous research on perceived objectivity

One of the novelties of the present research was to investigate whether people's intuitions about moral objectivity might be multi-dimensional rather than one-dimensional, as research concerning perceived objectivity (Goodwin and Darley 2008, 2012; Wright, Cullum & Schwab 2008; Wright, McWhite & Grandjean 2014), and research on moral diversity and tolerance (Haidt, Rosenberg, and Hom; Skitka et al. 2005) has thus far assumed. People's responses to a range of distinct views were tested and the question arose whether these conceptually distinct views are also psychologically distinct. The studies reveal that they do not directly map onto how people think about morality. As it turns out, relativism and no-truth form a single dimension. Similarly, universalism and absolutism are predicted by a shared psychological construct. Divine command theory, on the other hand, does form an independent psychological dimension. What do these results teach us about how people think about morality?

The crucial finding is that people conceive of the objectivity of moral judgments in different ways. There are some who take moral objectivity to imply that moral judgments are true independently of particular groups or cultures. There are others who take it to imply that moral judgments are true because they are based on universal and absolute moral norms, and yet others because moral judgments depend on divine commands. Each of the dimensions of Independent Truth, Universal Truth, and Divine Truth fulfills a distinct explanatory role. The validity of the FMO-scale that accommodates these three dimensions was established relying on previous research concerning perceived objectivity and its relation to tolerance. At the same time, the studies reveal that perceived objectivity is only one aspect of folk moral objectivity. A substantial amount of the variance in moral thinking can be

explained in terms of Independent Truth, Universal Truth and Divine Truth, which suggests that these constructs play an explanatory role in people's moral thought.

Studies 2 and 3 reveal that perceived objectivity scores require careful interpretation. Above I considered what high scores on perceived objectivity might mean. It is also worthwhile to explicate how low scores could be interpreted. People who do not perceive morality as objective might conceive of moral truth as being relative, or not believe in moral truths at all (if they score low on the dimension of Independent Truth). Alternatively, they might reject the idea of absolute moral norms that are universally binding (if they score low on Universal Truth). Finally, they might deny that there are true moral judgments that are based on divine commandments. The present results also shed a new light on research that shows a relationship between religion or divine command ethics and moral objectivity (Goodwin and Darley 2008; Sarkisian and Phelan 2019; Yilmaz & Bağcıkapılı 2015). Studies 2 and 3 replicate this finding in that Divine Truth is positively correlated with perceived objectivity. However, this relationship disappears when other dimensions are taken into account. As it turns out, religious grounding itself does not implicate an increase in perceived objectivity.

Finally, studies 4 and 5 reveal interesting correlations between the three dimensions on the one hand and tolerance and willingness to harm on the other. Study 3 shows that, rather than moral objectivity in general, scores on Universal Truth in particular are correlated with how comfortable people are to have someone with different moral views as a long-term guest. Study 5 shows that this relationship also applies to a measure of willingness to harm others. The study shows that high scores on Independent Truth strongly decrease people's willingness to harm others. This reveals that it can make a difference for which reasons people subscribe to moral objectivity. It is not the case that folk attitudes toward moral truth as such explain willingness to harm. Instead, one kind of moral objectivity is positively associated with it, whereas another is negatively associated with it. An important question for

further research is why this holds and whether it extends to actual forms of moral behavior.

9.2 *Relation to other research in social psychology*

Above I remarked that research on the psychological distinction between moral and non-moral attitudes makes a variety of assumptions that are relevant to the present research. For example, some research assumes that people generalize moral rules and violations to other social contexts while conventional rules and violations are perceived as applying locally (Turiel, 1983, 1998; Smetana, 1981, 1983; Smetana and Braeges, 1990). There is also research that shows that people are least supportive of moral diversity compared to other kinds of diversity (Haidt, Rosenberg & Hom 2003). Additionally, in research by Skitka and colleagues (e.g. Skitka et al. 2005) it is assumed that moral issues are perceived as objective, absolute, and universal, and their research shows that strong moral convictions, compared to strong non-moral attitudes, explains unique variance in interpersonal tolerance. The present research suggests that what is special about moral attitudes, compared to non-moral attitudes, varies on at least three distinct dimensions. While Study 4 shows that it is specifically Universal truth that predicts interpersonal tolerance, Study 5 shows that Independent Truth and Universal Truth even pull into different psychological directions when it comes to willingness to harm. Future research on moral attitudes and moral behavior might be advanced by taking into account the distinct role of each of the dimensions.

10. Conclusion

Existing research measures folk attitudes towards moral truth in terms of perceived objectivity. The findings presented here show that people have more fine-grained intuitions about moral objectivity. They can be captured in terms of three constructs, to wit Independent Truth, Universal Truth and Divine Truth. Each of these constructs is differently related to perceived objectivity, tolerance, and a willingness to harm measure.

This also provides additional insight into what distinguishes moral attitudes from non-moral attitudes. Existing experimental research in experimental philosophy and social psychology tacitly assumed that philosophical views such as absolutism, universalism, the rejection of relativism, or the idea that there are no moral truths, all exist on a single dimension. The present research shows that those views do not exist on a single dimension and that they do not function as a psychological conjoint, reinforcing each other. Rather, they play different psychological roles and they are different related to perceived objectivity, tolerance, and people's willingness to punish those who violate moral norms. In addition to providing new insights, this research also provides a methodological tool, namely the FMO scale, for conducting experimental research to folk moral objectivity and to interpersonal tolerance.

3.

ARE PEOPLE IMPLICITLY MORAL OBJECTIVISTS?

Abstract

Many philosophers assume that people are moral objectivists. It is well known, however, that people sometimes explicitly deny moral objectivism and verbally report that moral truths are relative or that there are no true moral judgments. Moreover, recent psychological research shows large variation in folk moral objectivity. Contrary to what philosophers assume, it is therefore possible that people are not moral objectivists after all. That depends, however, on the question of whether existing psychological research measures the relevant commitments. In this paper, I argue that there are at least two ways in which people can be moral objectivists, namely implicitly and explicitly. It is possible to explicitly deny being a moral objectivist while being implicitly committed to it. Enoch (2014) presents three tests to convince his reader that they are moral objectivists even if they explicitly think otherwise. As it happens, Enoch's tests, which he uses as intuition pumps, provide excellent measures of implicit metaethical commitments. In this paper I therefore use them as material for a survey study to test whether people are implicit moral objectivists. Results show that a significant majority of people is inclined toward moral objectivism and a minority is not. Consequently, despite people sometimes denying moral objectivism and the large variation found in some psychological studies, philosophers might be correct that most people are moral objectivists.

Keywords: Moral objectivism, folk metaethics, implicit, explicit, survey

1. Are People Implicitly Moral Objectivists?

Many philosophers think that people are moral objectivists (Cuneo 2007, Enoch 2014, Joyce 2006, Mackie 1977, Smith 1994). Yet, existing research suggests that they are in fact divided on this issue. There are large differences in objectivity ascriptions both between individuals and between different moral statements (Goodwin & Darley 2008, 2012, Wright, Cullum & Schwab 2008), which has been termed metaethical pluralism (Wright, Grandjean & McWhite 2013). And Sarkissian et al. (2011) show that people might be moral relativists. Furthermore, research by Pölzler and Wright (unpublished) suggests, while testing whether people's responses fit a multitude of realist and antirealist views, that people are mostly antirealists.

It might be, however, that existing studies do not measure people's real allegiances. A first reason is that existing studies on folk metaethics have suffered from construct validity (Pölzler 2018). The measurement instruments used did not always, or not exclusively, measure moral objectivism. The recommendations that Pölzler (2018) provides are really helpful in improving research on folk metaethics but there may be yet a larger issue that he did not account for. That is, most psychological studies asked questions that were more or less of an explicit metaethical nature. The nature of those questions may have prompted people to provide explicit metaethical beliefs and/or to theoretically speculate about the nature of morality. Yet, philosophers do not base the assumption that people are moral objectivists on explicit metaethical beliefs or theoretical speculations but on commitments or beliefs that are more implicit.

For example, Brink writes:

My appeal to commonsense moral thinking is not a prediction about the likely results of a Gallup poll on the issue of moral realism. Rather, my concern is with the philosophical implications or presuppositions of moral thought and practice. (...) I claim that cognitivism seems to be presupposed by common normative practices of moral judgment, argument,

and deliberation and that reflection on the nature of moral theorizing seems to support a realist view about these moral facts and truths. This claim may be false, but this is not shown by an appeal to common metaethical beliefs (or the lack thereof). (Brink 1989, p. 25)

In a later passage, Brink (1989, p. 51) literally refers to “[t]he objectivism or realism that is implicit in commonsense morality” (*my italics*).

Björnsson (2012, p. 9) draws an analogous distinction between explicitly commenting on and being engaged in moral thinking and debate. He goes on to argue that “the primary task of metaethical theories is to account for this engaged behavior, rather than for what is in effect lay people’s theoretical interpretations of it.” (*ibid.*: 9) This suggests that explicit beliefs need not be a reliable guide to people’s commitments concerning moral objectivity.

Similarly, Enoch (2005, p. 773, footnote 31) writes that:

[W]hat is relevant is not the explicit metanormative beliefs – much less the explicit metanormative statements – of participants in normative discourse. What is relevant, rather, are the deep metanormative commitments embedded (perhaps implicitly) in normative discourse and practice themselves. The fact that many sophomores (and not only them) express some subjectivist or relativist metanormative intuitions thus has very little weight in assessing the commitments of normative discourse.

The point can also be made in terms of Enoch’s (2014) distinction between what people explicitly think or say versus what they are implicitly committed to. Enoch argues that explicit commitments can diverge from implicit commitments. He goes as far as claiming: ‘You may think that you’re a moral relativist or subjectivist – many people today seem to. But I don’t think you are’ (Enoch, 2014, p. 193).

Existing psychological studies may have mostly measured people’s explicit commitments and, if so, this leaves open the possibility that people are committed to moral objectivism after all. To wit, existing studies ask participants whether or not particular moral claims can be true or false (‘truth-aptness task’) and/or whether it is possible that in an apparent moral disagreement one of two parties must be wrong or that both can be correct (‘disagreement task’). The metaethical nature of these questions may have induced people to engage in explicit theorizing about the status of morality. Moreover, the fact that they are asked to evaluate a range of moral and non-moral sentences as part of their participation in those studies increases the likelihood that they actively speculate about the answers to these questions. Indeed, it decreases the likelihood that they provide answers that reflect their implicit commitments.

I do not claim that existing studies solely measured people’s explicit commitments or that they lacked in sophistication. Indeed, existing studies used sophisticated measures and it is possible that those measures at least partly measured people’s implicit commitments. However, we cannot rule out that the findings of existing studies reflect people’s explicit commitments to such an extent that, if people are implicit moral objectivists, we fail to observe that this is the case. If so, existing studies leave open the possibility that people would, under different circumstances, unwittingly respond as moral objectivists.

In order to determine whether this is the case, people’s implicit commitments should be elicited. If it is true that people are moral objectivists, we should observe in measurements that elicit more implicit commitments that people are inclined toward moral objectivism. Fortunately, Enoch has developed three intuition pumps that are developed to show that people are moral objectivists even if they explicitly deny this. In this paper, I present the results of a survey experiment that features these three tests.

1.1 Enoch’s Tests

My use of the term “explicit commitments” is based on the descriptions provided above (e.g. Brink 1989, Björnsson 2012, Enoch 2014). It seems plausible

that there is a distinction between what people theoretically believe is the case, once they consider things explicitly, and what they are implicitly committed to when engaged in moral language and thought. I will therefore follow Enoch (2014) in the sense that there is a difference between implicit and explicit commitments and that when it comes to FMO, implicit and explicit commitments may diverge.

The goal of Enoch's tests is to make the reader realize that s/he is a moral objectivist, even if s/he explicitly denies being one. In other words, they serve to lay bare people's implicit commitments. Enoch presents them as thought experiments or intuition pumps. The underlying idea is that when confronted with adequate stimuli it should take only a moment of reflection for people to realize that they implicitly accept the objectivity of morality. Or, at least, that is Enoch's conjecture. But is that indeed the case? Do people, when confronted with Enoch's tests, indeed respond in a way that shows that they are moral objectivists? This question is important because it has been assumed that people's objectivist commitments provide *prima facie* support for the truth of moral objectivism and for different varieties of moral realism. Moreover, the question is important because existing psychological research shows a large variety in objectivity ascriptions both between individuals and between different moral issues. A study on implicit commitments, as presented here, may provide different results. In order to answer this question, I use Enoch's tests as a means to empirically examine whether people are indeed implicitly committed to moral objectivism.

2.1 *Test 1: Jokes*

Consider the following joke:

[TASTE] A child hates spinach. He then responds that he's glad he hates spinach. To the question "Why?" he responds: "Because if I liked it, I would have eaten it; and it's yucky!" (Enoch, 2014, p. 193)

What makes this joke funny, according to Enoch, is that the child fails to realize that yuckiness is a subjective property. Given that he imagines that his likes and dislikes change, he should adjust his judgment about the taste of spinach. But he does not do this. In other words, the child misconstrues yuckiness as being independent of his likes and dislikes. The child treats yuckiness as an objective property while we, as the reader, are inclined to treat it as a subjective property. That is why the story is funny.

The joke would not work, Enoch proposes, if the subject matter is something that we deem to be more objective than taste. To make his point, Enoch presents a story about factual matters that has the same structure:

[FACTUAL] Consider, for instance, someone who grew up in the twentieth-century West, and who believes that the earth revolves around the sun. Also, she reports to be happy that she wasn't born in the Middle Ages, "because had I grown up in the Middle Ages, I would have believed that the earth is in the center of the universe, and that belief is false!" (Enoch, 2014, p. 193-194)

Clearly, this version of the joke does not work. The reasoning of this person sounds perfectly sensible while the reasoning of the child does not. Enoch explains:

[i]f the joke works, this seems to indicate that the subject matter is all about us and our responses, our likings and dislikings, our preferences and so on. If the joke doesn't work, the subject matter is much more objective than that, as in the astronomy case. (Enoch, 2014, p. 193-194)

Given that we are interested in morality, the question then arises whether moral versions of the joke more closely resemble taste (and hence, are deemed to be about subjective matters) or [factual] (and hence, are deemed to be more objective). Enoch thinks that we, philosophers and lay people alike,

are implicitly committed to regarding morality as more similar to matters of fact as compared to matters of taste. To pump this intuition, Enoch provides a third version of the joke:

[MORAL] Suppose someone grew up in the US in the late twentieth century, and rejects any form of racism as morally wrong. He then reports that he's happy that that's when and where he grew up, because "had I grown up in the 18th century, I would have accepted slavery and racism. And these things are wrong!" (Enoch, 2014, p. 194)

As with the factual case, this sounds perfectly sensible and the joke is clearly not funny. Enoch believes that this is best explained by the fact that we tend to ascribe objectivity to morality.

2.2. Test 2: *The phenomenology of disagreement*

The second test that Enoch presents concerns the phenomenology of disagreement, or what it feels like from the inside when people have a difference of opinion. Enoch claims that different types of disagreements have a different feel. When people disagree about whether human actions influence global warming, this feels like trying to get an objective fact right. In contrast, a disagreement about whether bitter chocolate tastes better than milk chocolate feels like stating one's own preferences. The test consists of answering the following question: Does a moral disagreement feel more like a disagreement about whether bitter chocolate tastes better than milk chocolate or does a moral disagreement feel more like a disagreement about whether human actions influence global warming? Enoch conjectures that a moral disagreement is more similar to a disagreement about global warming as compared to the one about bitter chocolate. The explanation that Enoch provides for this is that a moral disagreement feels more like a matter of getting the objective facts right.

2.3. Test 3: *The counterfactual question*

Enoch's third test concerns counterfactuals. The key question is: "Had our beliefs and practices been very different, would it still have been true that so-and-so?" (Enoch 2014, p. 197) Consider the following question: Had we believed that smoking is harmless and did not ban smoking, would it still have been true that smoking causes cancer? Clearly, we would answer affirmatively. Enoch suggests that this is the case because we believe that it is an objective fact that smoking causes cancer. He compares this example with our beliefs and practices concerning gender-discrimination. Had we believed that gender-discrimination is ok, would it then still have been wrong? Enoch argues that we would also say "Yes" in this case. Enoch thinks that we treat our moral beliefs as objective facts, and moral truth as independent from our beliefs and practices, and therefore the answer must be "Yes".

Enoch's three tests are very clever and his explanation of how we would respond seems plausible. If people are indeed moral objectivists, the test should show that this is the case. Importantly, the tests do not concern people's beliefs about morality. Instead, their answers are meant to reveal their implicit commitments about it. For instance, if you do not think that [moral] is funny, then, Enoch proposes, you must at some level regard morality as more objective than taste. The same is true if you take the phenomenology of moral disagreement to resemble factual disagreement, and if you take what is wrong to be invariant with respect to our beliefs. If you are a moral objectivist, in an implicit sense, then you should respond as Enoch predicts. However, when used as intuition pumps, these tests are of limited value. It can hardly be taken for granted that the responses of a particular individual generalize. Because of this, there is ample reason to put this to the test in the form of a survey experiment.

3. Experiment

In order to test whether ordinary people are implicitly committed to moral objectivism, I used Enoch's tests of moral objectivity as a basis for an experiment.

Participants first read each of the stories (taste, factual, moral) at random and, directly following each story, were asked to indicate whether the story could be perceived as a joke and to rate the degree to which they believed the story was funny. Secondly, participants were asked to indicate how moral disagreements “feel”. Do they feel more like disagreements about bitter chocolate versus milk chocolate or do they feel more like disagreements about whether human actions contribute to global warming? Thirdly, participants were asked whether certain actions that they now deem morally wrong would still be morally wrong had our beliefs and practices been different. To be certain that participants’ read the instructions carefully and understood the questions, they were also presented with two checks to test whether they interpreted the questions in the way required. Taken together, the results of this experiment should provide us with insight into the question of whether or not people are moral objectivists in an implicit sense.

3.1. *Methods*

3.1.1. *Participants*

150 participants living in the United States were recruited via the online service Mechanical Turk and received \$0.75 for their time. There were some participants who did not complete the survey and a few who failed to answer correctly two multiple-choice questions that assessed whether they understood the second and third test appropriately and those participants were excluded from statistical analyses.¹³ Analyses were conducted on the remaining 97 participants (45 female; *Mage* = 37). The checks followed after the questions about how moral disagreements feel and the question whether or not

actions are still wrong when beliefs and practices change. In the final part, participants answered three demographics questions (age, gender, nationality).

3.1.2. *Materials and Procedure*

The study consisted of three different parts. In the first part, participants received the three different versions of the Enoch jokes (taste, factual, moral) in a random order. In each case, they were asked to answer two questions that serve to measure their implicit commitments: (1) “Can the above story be regarded as a joke?” (Yes, No); (2) “To what extent do you think the above story is funny?” (0-100).

In the second part, participants read a very short story based on Enoch’s explanation about the difference between disagreements about taste and disagreements about factual matters (Enoch 2014, p. 195-196). Participants were subsequently asked to think about a moral disagreement about abortion or a disagreement about a different moral issue they felt strongly about. They were then presented with the question whether the moral disagreement they have in mind feels more like a disagreement about whether bitter chocolate tastes better than milk chocolate or whether it feels more like a disagreement about whether human actions contribute to global warming (for materials, see Appendix 3).

In the third part, participants read a short story based on Enoch’s explanation about how smoking would still cause cancer had our beliefs and practices changed (Enoch 2014, p. 196-197). They were subsequently introduced with a counterfactual of the form “Had our beliefs and practices been very different, would it still have been true that so-and-so?” They were asked whether gender-discrimination (or a different moral issue they feel strongly about) would still be wrong had our relevant practices and beliefs been different (for materials, see Appendix 4).

¹³ More information on the attention checks can be found in Appendix 1 at the end of the paper. There were no differences in results between participants who completed the survey and answered the attention checks correctly and participants who completed the survey but failed on the attention checks. Moreover, I have conducted this exact survey experiment twice on two different data samples. The first study did not contain attention checks and, to be certain of valid results (i.e. Polzler, 2018), I therefore conducted the same study a second time including attention checks (four months after the first data collection). In Appendix 2, I provide the results of the data that I collected for the study without attention checks. The fact that the results of both survey experiments are very similar suggests that the effects found here are fairly robust.

3.1.3. Results

Table 1. Mean Proportions and 95% confidence interval for the question “Can the above story be regarded as a joke?”

	Spinach	Factual	Moral
Yes	63.9 ± 9.55	32 ± 9.29	8.20 ± 5.46

Figure 1. Differences between the Taste, Factual, and Moral categories

Table 1 presents the percentages of “Yes” responses for each of the versions of the jokes and the upper and lower boundaries of the 95% confidence interval. Figure 1 shows that the 95% confidence intervals do not overlap and hence that the differences between each condition are significant. On the basis of the 95% confidence intervals, we can estimate the population mean for the proportion of “Yes” responses for the question of whether the taste story can be regarded as a joke. This lies between 54.34% and 73.46%. For the factual story, the population mean lies between 22.72% and 41.28%, and for the moral story, the population mean lies between 2.74% and 13.66%.

Because of the within-subjects design of the present study and the non-normal distribution of scores, a non-parametric test, the Wilcoxon Rank Sum Test, was conducted to test for differences in perceptions of how funny the joke is. The results indicate that the taste story was perceived as funnier (mean rank = 44.36) than the factual story ($Z(94) = -3.551, p < .001$). Results also show that the taste story was perceived as funnier (mean rank = 44.29) than the moral version ($Z(94) = -6.676, p < .001$). The factual story was also perceived as funnier (mean rank = 37.47) than the moral story ($Z(94) = -5.531, p < .001$).

Figure 2. Proportions for the question about the phenomenology of disagreement

Figure 3. Proportions for the question about the counterfactual “Had our beliefs and practices been different...”

Proportion tests were conducted both for the question about how a disagreement feels and for the question about the counterfactual of the form “Had our beliefs and practices been different, would it still have been wrong that so-and-so?” (see figure 2 and 3). For the disagreement question, there were four participants who selected “Other”. These participants indicated that moral disagreements do not feel like disagreements about taste nor like factual disagreements. For example, one participant remarked, “It feels different. Moral issues are not a matter of taste nor an objective truth”. Another participant remarked, “I think it’s a little bit of both. There are some objective truths to morality but also some subjective opinions”. Proportions were calculated for the remaining participants ($N = 93$). Of those participants, 22.5% responded that moral disagreements feel like disagreements about matters of taste and 77.5% that they feel like factual disagreements.

For each question a proportion test was conducted with the null hypothesis that the proportion is 0.5 versus the alternative hypothesis that the true population proportion is higher than 0.5. For the disagreement question, results show that the null hypothesis is rejected ($Z(93) = 5.41, p < .001$). This implies that participants were significantly more inclined to answer that moral disagreements feel like factual disagreements compared to disagreements about taste.

For the counterfactual question, a proportion test similarly showed that the null hypothesis is rejected ($Z(97) = 3.49, p < .001$). At the same time, 22.5% of participants indicated that a moral disagreement feels like a disagreement about taste and 29.9% of participants indicated that had our beliefs and practices been different then abortion (or a different moral issue they felt strongly about) would not be morally wrong. After combining the proportions of the disagreement and counterfactual test, it turns out that at most 57.7% of participants provide objective responses.

2.2. Discussion

Results of the experiment show that most people support moral objectivism in an implicit sense. With regard to the first test, people perceive the taste story as a joke but not the factual and moral versions. Moreover, most people perceive the taste story as funnier than the factual or moral story. Accordingly, the first test suggests that people treat yuckiness as a subjective property and moral and factual matters as involving non-subjective properties. Interestingly, the difference between the factual story and the moral story is statistically significant. This suggests that people treat factual and moral stories differently. It is possible, however, that participants exaggerated a bit because of the sensitivity of moral issues.

Results also show that for a majority of people moral disagreements feel more like factual disagreements than like disagreements about taste. Following Enoch’s reasoning, the second test suggests that the feel of moral disagreements is more about getting an objective fact right than about stating one’s own preferences. Similarly, the third test shows that a majority of people believes that moral wrongness is fixed over different counterfactual circumstances. That is, had their beliefs and practices been different, they would still judge gender-discrimination as morally wrong. Overall, these results provide support for Enoch’s thesis that people are moral objectivists.

In light of recent studies providing mixed evidence on the idea that people are moral objectivists (Goodwin and Darley 2008; 2012; Wright et al. 2013), or that appear to show that people are downright antirealists (Pölzler and Wright, unpublished), this study represents a departure from that pattern. The main difference with previous studies is that the present research deliberately tests people’s implicit commitments.

The different pattern observed here supports the idea that it makes sense to distinguish between different types of metaethical commitments. Existing research on folk moral objectivity may have been tapping into the wrong type of commitments, namely explicit commitments, or results may have been conflated by reflecting both implicit and explicit commitments, while the

objectivity that philosophers assume to be characteristic of ordinary moral discourse and practices is much more implicit. If so, these philosophers may have been right after all¹⁴.

Taking into consideration the distinct pattern of responses found here compared to previous research, it may be helpful for future psychological research to distinguish between implicit and explicit commitments. Similarly, for philosophers who believe that people are (either implicitly or explicitly) moral objectivists, and want to make this assumption amenable for empirical testing, the general recommendation is to be precise about how specific commitments reveal themselves in ordinary moral discourse and practices. More specifically, the recommendation is to make clear whether they are implicit and/or explicit commitments. Because we need more evidence, the findings here do not conclusively show that people are implicit moral objectivists. However, the results do point in that direction.

4. Conclusion

I have presented a survey study concerning the question whether people are moral objectivists. I used three tests that are designed to get at people's implicit commitments about moral objectivity. The study reveals that a significant majority of people regards morality as more objective than taste and as similar to matters of fact. I therefore submit that, in spite of existing research suggesting that people may be metaethical pluralists or metaethical antirealists, this study provides some support for the idea that people might be moral objectivists after all. That is, implicitly.

¹⁴ What should be noted is that there is a sizable minority of respondents providing a response that contradicts FMO. One possibility is that FMO appeals to some people but not all. A different possibility is that the tests developed by Enoch may not be ideal for measurement instruments. Yet another possibility is that there are large differences in metaethical intuitions between different individuals and that some examples elicit objectivist intuitions in some individuals but not in others. We can only find an answer to this question by testing these possibilities in future research. At this time, I interpret the finding that a significant majority of people respond as suggested by Enoch as defeasible evidence in favor of FMO.

Appendix 1: Attention Checks

In what follows we have two comprehension questions that we would like you to answer.

In part II of this survey we asked you about how it feels for you to engage in a certain kind of disagreement. Please select one of the questions below that, in your view, captures best what we were aiming at in part II.

- [1] Is abortion morally permissible?
- [2] Does a moral disagreement feel more like a disagreement about subjective preferences or more like a disagreement about getting the objective facts right?
- [3] Are disagreements about chocolate the same as disagreements about global warming?
- [4] I do not know.

In part III of this survey we asked you a question whether gender-based discrimination (or a different example, if you think it is permissible) would still be wrong if our relevant practices and beliefs were different. What was the aim of that question?

- [1] To find out whether you believe that gender-based discrimination is morally wrong.
- [2] To find out whether you believe that smoking still causes cancer if our relevant practices and beliefs concerning smoking would be different.
- [3] To find out whether you believe that moral claims are, similar as scientific claims such as "smoking causes cancer" about objective matter of fact.
- [4] I do not know.

Appendix 2: Results first survey experiment (see footnote 1)

Table 1 Mean proportions and 95% confidence interval for the question “Can this be regarded as a joke?”

	Spinach	Science	Moral
Yes	59.6 ± 10.12	26.6 ± 9.11	11.7 ± 6.21

Table 1 presents the percentages of “Yes” responses for each of the versions of the jokes and the upper and lower boundaries of the 95% confidence interval for the first survey experiment (without attention checks). These results show that the spinach version is perceived considerably more often as a joke than the science and moral versions. The population proportion of “Yes” responses for the question of whether the spinach version of the story can be regarded as a joke lies between 49.48% and 69.72%. For the science version the population mean lies between 17.49% and 35.71%, and for the moral version it lies between 5.49% and 17.91%.

Because of the within-subjects design of the present study and the non-normal distribution of scores, a non-parametric test, namely the Wilcoxon Rank Sum Test, was conducted to test for differences in how funny participants found each of the jokes. The results indicate that the spinach version was perceived as funnier (mean rank = 37.24) than the science version ($Z(97) = -4.042, p < .001$). Results also show that the spinach version was perceived as funnier (mean rank = 40.75) than the moral version ($Z(97) = -6.893, p < .001$).

Proportion tests were conducted both for the question about how a disagreement feels and for the question about counterfactual reasoning. For the former question, a proportion test was conducted with the null hypothesis that the proportion is 0.5 versus the alternative hypothesis that the true population proportion is higher than 0.5. The results show that the null hypothesis can be rejected ($Z(97) = 3.92, p < .001$). This implies that participants were more inclined to answer that moral disagreements feel like factual disagreements. For the question about counterfactual reasoning, a proportion test similarly showed that the null hypothesis can be rejected ($Z(97) = 3.50972, p < .001$). At the same time, 26.6% of participants indicated that a moral disagreement feels like a disagreement about taste and 30.9% of participants indicated that had our beliefs and practices been different, then abortion (or a different moral issue they felt strongly about) would not be morally wrong. Taken together (i.e. combining scores of the phenomenology of disagreement test and the counterfactual test), 53.6% of participants responded that moral disagreements feel like disagreements about global warming and that moral issues would not be wrong had beliefs and practices been different.

Appendix 3: The phenomenology of disagreement

People engage in all sorts of disagreements. For example, we may engage in a disagreement about whether bitter chocolate tastes better than milk chocolate. We can also disagree about whether human actions influence global warming. These two disagreements are, however, different. In the chocolate case, it feels like stating one's own preference, and perhaps trying to influence the listener into getting his own preferences in line. In the global warming case, though, it feels like trying to get at an objective truth, one that is there anyway, independently of our beliefs and preferences. That is, either human actions contribute to global warming, or they do not.

In this part of the study, we will consider what it feels like for you to engage in a disagreement.

Now, think of some serious moral disagreement. For example, about the moral status of abortion. Suppose that you are engaged in such a disagreement. Imagine this, as it were, from the inside. You are in this disagreement yourself. Perhaps you think that there is nothing wrong with abortion, and you are arguing with someone who thinks that abortion is morally wrong. Or, perhaps you think that abortion is morally wrong and you are arguing with someone who thinks that there is nothing wrong with it.

Please explain how it feels for you to engage in this kind of disagreement. Please note that there is no correct answer to this question. We would simply like to know how it feels for you to engage in moral disagreements. In particular, please tell us whether it feels more like disagreeing over whether bitter chocolate tastes, or like disagreeing over objective facts like whether human actions contribute to global warming or not?

- [1] It feels more like disagreeing over which chocolate is better
- [2] It feels more like disagreeing over whether human actions contribute to global warming
- [3] Other

Appendix 4: The counterfactual question

As a result of years of scientific research we now know that smoking causes cancer. Now, had our relevant practices and beliefs regarding smoking been different - had we been ok with it, had we not banned it, had we thought smoking was actually quite harmless - would it still have been true that smoking causes cancer? It is probably uncontroversial that the answer is "Yes". The effects of smoking on our health do not depend on our beliefs and practices. Rather, it is an objective matter of fact.

The question that we therefore ask here is "Had our beliefs and practices been very different, would it still have been true that so-and-so?"

Let us apply this question to morality. For example, some people believe that gender-based discrimination is wrong. Maybe you also believe that it is morally wrong or maybe you do not. If you do not, imagine something else that you think is morally wrong. Would it still have been wrong had our relevant practices and beliefs been different?

- [1] No, had our relevant practices and beliefs been different than it would not be wrong.
- [2] Yes, had our relevant practices and beliefs been different than it would still be wrong.
- [3] Other

4.

DO PEOPLE BELIEVE IN OBJECTIVE MORAL PROGRESS, KNOWLEDGE, AND ERROR?¹⁵

¹⁵ This chapter is co-authored by Thomas Polzler and Jacob Dijkstra.

Abstract

A prevalent assumption in metaethics is that lay people believe that moral judgments are objectively true or false. This is the assumption of folk moral objectivism (FMO). FMO has been investigated by testing whether or not people believe that moral statements were true or false and whether moral disagreements imply that at least one party is mistaken. In this paper we argue that FMO also entails that people believe in the possibility of objective moral progress, moral knowledge, and moral mistakes. Moreover, we developed a survey to test whether this is indeed the case. We compare how people respond to statements from four different domains: preferences, conventions, science, and morality. And we test, both abstractly and concretely, whether or not in each of the domains people believe that objective progress, knowledge, and mistakes are possible. Our results show that people tend to believe that, in contrast to the scientific domain, this is not possible in the moral domain.

Keywords: Folk moral objectivism, Moral progress, Moral knowledge, Moral error, survey

1. Introduction

A central debate in metaethics concerns the existence of objective moral truths. These include in particular statements about whether states of affairs are good or bad and whether actions are right or wrong. In order for such statements to be objectively true or false, our judgments about such properties must be true or false independently from the perspective or the beliefs of specific individuals and/or cultures. Some philosophers affirm the existence of objective moral truths (Boyd 1988; Brink 1989, Smith 1994) while others deny it (Ayer 1936; Blackburn 1993; Mackie 1977). Most philosophers agree, however, that ordinary people are moral objectivists, the thesis of folk moral objectivism that I have discussed before (FMO). Indeed, a prevailing assumption is that metaethical theories should accommodate FMO.

There are even philosophers who argue that the fact that people are moral objectivists supports the thesis of moral objectivism itself, which is an argument that is called “the presumptive argument” (e.g., Brink 1989; Enoch 2017a, 2017b; Dancy 1986; Huemer 2005; McNaughton 1988). According to this argument, the fact that people believe that moral judgments are objectively true or false provides a *prima facie* reason to believe that moral objectivism is true. In other words, they claim that metaethical theories should take moral objectivism as point of departure unless there are strong arguments against the existence of objective moral truths. A particularly influential formulation of this argument can be found in Brink:

In many areas of dispute between realism and antirealism, realism is the natural metaphysical position. [...] So too, I think, in ethics. [...] if this claim about the realist nature of moral inquiry is right, we have reason to accept moral realism that can be overturned only if there are powerful objections to moral realism. (Brink 1989, p. 23-24)¹⁶

¹⁶ Theories in metaethics that are ‘realist’ endorse the idea that (at least some) moral judgments are objectively true or false.

The presumptive argument is not entirely undisputed (for instance, see Loeb 2007; Pölzler 2018) but it has not been scrutinized empirically. If FMO is false, the argument loses much of its force.

According to Brink (and other proponents of the presumptive argument), five features of ordinary moral thought and practice particularly show that people take morality to be objective and are indicators of FMO:

1. **INDEPENDENCE:** People believe that moral sentences are true or false independently of anyone's subjective reactions or attitudes.
2. **EXCLUSION:** People believe that only one party in a moral disagreement can be correct.
3. **PROGRESS:** People believe that it is possible to make moral progress.
4. **KNOWLEDGE:** People believe that it is possible to have moral knowledge.
5. **ERROR:** People believe that it is possible to make moral mistakes.

Independence and Exclusion have been investigated in previous research (i.e. Goodwin and Darley 2008, Wright et al. 2013, 2014; Pölzler 2018; Pölzler and Wright forthcoming) and we believe that it would also be valuable to investigate Progress, Knowledge, and Error. If we obtain empirical evidence for FMO then this would support the idea that FMO should be accommodated by metaethical theories. In contrast, if we obtain evidence that is in tension with FMO, we may want to consider to what extent metaethical theories should accommodate FMO. Either way, if we would obtain empirical results that pertain to FMO, we would potentially benefit metaethical theorizing and we would help philosophers to assess the plausibility of metaethical claims and theories.

Recent psychological research on folk metaethics has investigated Independence and Exclusion. *Prima facie*, the results of those studies do not univocally support FMO. There seems to be a large inter- and intrapersonal variance in how people think about the truth and falsity of moral judgments and whether or not moral disagreements imply that at most one party can be correct. However, in light of the discussion in Chapter 1, we should be careful in interpreting those results as evidence against FMO in light of issues of construct validity.

Moreover, some philosophers believe that FMO applies to implicit commitments (Brink, 1989; Björnsson 2012; Enoch, 2014) and as discussed in Chapter 1 and 3, it is not entirely clear whether existing empirical research provides evidence for implicit as opposed to explicit commitments. What we therefore need is measurement instruments that correctly identify what people are implicitly and explicitly committed to. We will then be able to answer the question whether people are implicitly and/or explicitly committed to FMO.

In the research that is part of this paper, we will employ measures that are rather explicit. The advantage of using explicit measurement instruments is that they provide us with the possibility to carefully distinguish different metaethical views in the experimental stimuli that we use. Of course, there is a valid worry that relatively explicit measures also elicit explicit commitments, which according to some philosophers is irrelevant for FMO.

However, the fact that measurement instruments are explicit does not imply that they necessarily only measure explicit commitments. Moreover, given that we use scenarios of previous research as input, this will provide us with the possibility to compare our results with those of existing research.

The question that is addressed here is whether people believe in moral progress, knowledge and error. To answer these questions, we conducted a survey experiment.

2. The present research

We investigate whether people believe in moral progress, knowledge, and error. Importantly, there is a difference between believing that moral progress, knowledge, and/or error is possible versus that it obtains. FMO requires that people believe in the possibility of progress/knowledge/error and not necessarily that moral progress has occurred, or that moral knowledge is obtained, or that moral mistakes are made. We therefore use measurement instruments that aim at testing whether people believe in the possibility of moral progress/knowledge/error.

Moreover, we would like to know whether people believe in moral progress/knowledge/error in a way that corresponds to FMO. It is possible, however, that people believe that moral progress, knowledge, and/or error is possible but not objectively speaking. For instance, perhaps people believe that there is only the possibility of progress/knowledge/error in a subjectivist sense. That is, according to the perspective or the beliefs of specific individuals and/or cultures. For instance, with regard to moral error, perhaps people believe that moral error is only possible in the sense of failing to grasp one's own moral beliefs or the beliefs of one's own culture.¹⁷ Of course it is also possible that people do not believe that moral judgments are objectively true or false and that moral progress, knowledge, or error is not possible.

Finally, existing research has shown that people tend to respond differently to abstract versus concrete cases (Nichols & Knobe 2007; Murray & Nahmias 2014). We therefore test people's responses both in an abstract and in a concrete condition. For these reasons, we developed a study in which people can indicate, both abstractly and for concrete cases, whether or not they think that moral progress/knowledge/error is objectively or subjectively possible or that it is not possible at all.

We hypothesize that FMO is true and that therefore people will, along the dimensions of progress, knowledge, and error, treat moral statements in ways similar as they treat scientific statements and different from how

¹⁷ Moreover, error theorists (according to whom all of our moral judgments are mistaken) also grant that moral errors are possible.

they treat statements about conventions and personal preferences. In statistical terms, our null hypothesis is that there are no differences, along the dimensions of progress, knowledge, and error, in how people treat moral statements versus other kinds of statements and our alternative hypothesis is that there are differences.

2.1 Method

2.1.1 Participants

We recruited 453 participants via the online service Amazon Mechanical Turk. They received \$0.80 for their time. To prevent that some participants do not take the cognitive effort needed to provide valid responses we included several attention checks and measured participants' survey completion times (see Huang et al. 2012; Pölzler forthcoming). 26 participants were excluded from statistical analyses because they failed with regard to at least one attention check. 18 participants were excluded because they finished the survey very fast. Analyses were conducted on the remaining 409 participants.

2.1.2 Materials and procedure

Participants were randomly distributed to an abstract version of the study or to a concrete version, making version a between-subjects factor. In each version participants were presented with a number of statements. Each statement belonged to one of four domains (morality, social conventions, personal preferences, and science) and was assessed along each of three dimensions (progress, knowledge and error). Dimension and domain were thus within-subjects factors. The abstract version for each domain consisted of the evaluation of a single abstract statement — addressing morality, social conventions, personal preferences and science in general — along all three dimensions. This implies a total of 12 evaluations. The concrete version involved three concrete statements for each of the social conventions, personal preferences, and science domains, and nine statements for the morality domain. This yielded a total of 18 concrete statements that were each evaluated on all

three dimensions, for a total of 54 evaluations. Each version (abstract and concrete) and each domain (morality, social conventions, personal preferences, and science) were randomized. The design is summarized in Table 1 below.

Table 1. Number of evaluations per version (abstract v. concrete) X dimension (progress, knowledge, error) X domain (morality, social conventions, personal preferences, science) combination.

Version	→	Progress		Knowledge		Error	
		Abstract	Concrete	Abstract	Concrete	Abstract	Concrete
Domain	morality	1	9	1	9	1	9
	conventions	1	3	1	3	1	3
	preferences	1	3	1	3	1	3
	science	1	3	1	3	1	3

Note: In the abstract version, four statements were each evaluated along the three dimensions; in the concrete version, 18 statements were each evaluated along the three dimensions

Measures

To determine whether a participant affirmed the possibility of objective moral progress, objective moral knowledge, and objective moral error, we used the following measures (with “X” being either filled by one of our abstract or concrete statements).

PROGRESS:

Does it seem to you that, at least over long periods of time, progress can be or has been made with regard to the question of whether X (i.e., answers to this question cannot only change but become better)?

- (1) There can be or has been progress in the sense of (coming closer to) discovering the objective truth about this question.

- (2) There can be or has been progress but only from the perspective of particular cultures or individuals.
- (3) There cannot be progress.

KNOWLEDGE:

Does it seem to you that the question of whether X is such that, under favorable circumstances, a person can know the answer to this question (i.e., can have a justified and true belief about it)?

- (1) It is possible to acquire such knowledge by investigating the objective truth about this question.
- (2) It is possible to acquire such knowledge but only by investigating one’s culture’s or one’s own beliefs about this question.
- (3) It is not possible to acquire knowledge about this question.

ERROR:

Does it seem to you that the question of whether physical punishment is morally wrong is such that, at least under favorable circumstances, a person can err about the answer to this question (i.e., can give a false answer to it)?

- (1) It is possible to make such an error in the sense of failing to grasp the objective truth about this question.
- (2) It is possible to make such an error but only in the sense of failing to grasp one’s culture’s or one’s own beliefs about this question.
- (3) It is not possible to make an error about this question.

Based on our considerations in Section 1, we interpreted (1) answers as indicative of FMO and (2) and (3) answers as indicative of responses that are in tension with FMO. Our hypotheses entail that our participants would predominantly be drawn towards (1) answers.

2.1.3 *Statements*

Although our hypotheses pertain to the moral domain, statements from non-moral domains were included for three reasons. First, to provide a benchmark against which to assess whether or not moral statements are interpreted in an objectivist sense. We assume that people will be drawn towards objectivity for scientific statements and towards subjectivity or non-objectivity for the statements about conventions and preferences. Second, the differences between these domains would prevent participants from answering all tasks in the same way. That is, the comparison of statements from different kinds of domains may increase people's ability to assess its metaethical grounding. Third, existing empirical results (Wright, Grandjean & McWhite 2013; Pölzler and Wright forthcoming) show that people make different distinctions between moral and non-moral statements (more on this below).

Here are the statements from the concrete task (organized by domain) that we used (some of which are taken from or based on Goodwin and Darley 2008, and Pölzler and Wright forthcoming).

MORALITY

Abortion is morally permissible.

Physical punishment is morally wrong.

It is good to do unto others as you would have them do unto you.

A country with the death penalty is morally worse than a country without.

Eating factory-farmed meat is morally bad.

Siblings ought not to kiss each other on the mouth passionately.

Selling children on the internet is morally wrong.

Cleaning one's bathroom with the American flag is morally impermissible.

Helping terminally ill patients end their lives is morally permissible.

SOCIAL CONVENTIONS

Wearing pajamas and bath robes to a seminar meeting is wrong behavior.

One ought not speak with one's mouth full.

Talking loudly and constantly to the person next to you during a lecture is a permissible action.

PERSONAL PREFERENCES

Getting tattoos and/or body piercings is okay.

Shakespeare is a better writer than is Dan Brown (author of *The Da Vinci Code*).
Classical music is the best kind of music.

SCIENCE

The earth is flat.

Boston (Massachusetts) is farther north than Miami (Florida).

The chemical formula for water is H_2O .

After completing all tasks described above, participants were again presented with the list of our 18 concrete statements and were asked the following request: "Below you find a number of statements. For each of these statements please indicate whether you think that it is primarily about morality, social conventions, personal preferences or scientific facts." This task allowed us to address the potential objection that participants' moral progress, knowledge and error responses are not explained by their metaethical commitments in favor of or against moral objectivity, but by their distinguishing moral from non-moral statements differently from how we did (Wright et al. 2013).

2.3 *Results*

2.3.1 *Statistical hypotheses and analytical strategy*

Within each dimension (progress, knowledge, and error) we compare answers to statements from the morality domain to answers from statements to the other domains. The statistical null hypotheses that we test are thus that

people are equally objectivist in the morality domain, along the dimensions of progress, knowledge, and error, compared to the other domains. Moreover, data from the abstract and concrete tasks are analyzed separately. Referring to Table 1 above, this implies that analyses are performed down each of the 6 columns, separately. The first two columns pertain to the abstract and concrete versions of measuring moral progress, which we will refer to as Hypothesis 1. The third and fourth columns pertain to the abstract and concrete versions of moral knowledge, which we will refer to as Hypothesis 2. Finally, the fifth and sixth columns pertain to the abstract and concrete versions of moral error, which will be referred to as Hypothesis 3.

For each statement participants indicated whether they believe objective progress, knowledge, or error to be possible (1) or not (2 or 3). For testing the hypotheses we dichotomize these answer categories into “objective” (1) and “not objective” (2 or 3). We employed a within-subjects design and observations (evaluations of statements) are therefore nested in participants, leading to dependent observations (e.g., Gelman and Hill 2007; Snijders and Bosker 2011). We account for this dependence by estimating multilevel logistic regression models with observations nested in individuals, using the Gibbs sampler implemented in the OpenBUGS language run from R (Gelman and Hill 2007, Lunn et al. 2009).

2.3.2 Descriptives

Overall, there are 409 participants, of whom 216 were assigned to the abstract version and 193 to the concrete version. Of the participants taking the abstract version 117 self-identified as male and 99 as female; mean age in the abstract version was 35.86 years (sd = 10.39). Of the participants taking the concrete version 109 self-identified as male, 82 as female, and 2 as “other”; mean age in the concrete version was 36.95 years (sd = 11.19). Table 2 below presents the percentages of “objective” answers for all statements, broken down by versions, dimensions, and domains.

Table 2. Overall percentages of “objective” answers to statements, by version (abstract v. concrete) X dimension (progress, knowledge, error) X domain (morality, social conventions, personal preferences, science) combination.

Version	→	Dimension					
		Progress		Knowledge		Error	
		Abstract	Concrete	Abstract	Concrete	Abstract	Concrete
Domain	Morality	46.76 ^a	33.79 ^c	48.15 ^a	32.53 ^c	45.37 ^a	27.58 ^c
	Conventions	47.69 ^a	29.53 ^b	54.17 ^a	33.33 ^b	51.39 ^a	28.15 ^b
	Preferences	49.54 ^a	25.04 ^b	50.46 ^a	22.80 ^b	43.06 ^a	22.80 ^b
	Science	84.26 ^a	53.71 ^b	83.33 ^a	68.39 ^b	78.24 ^a	46.29 ^b

Note: ^a based on 1 statement per participant; ^b based on 3 statements per participant; ^c based on 9 statements per participant; N = 216 participants in the abstract version and N = 193 participants in the concrete version

Table 2 reveals a very consistent pattern in the data, seemingly going against the null hypotheses. In both abstract and concrete tasks, and along all three dimensions (progress, knowledge, and error), scientific statements are consistently evaluated in a more objectivist way than statements from the other three domains. Importantly, the evaluation of moral statements does not seem to differ from the evaluation of statements from the conventions and preference domains.

To formally test the hypotheses, we estimate multilevel logistic regressions using the Gibbs sampler embedded in the OpenBUGS language. Evaluations of statements (level 1) are nested in participants (level 2). The response variable in all analyses is the logarithm of the odds of evaluating a statement with the answer “objective”. The model equation has a random term (intercept) for each participant. Uninformative priors are used for all (hyper-)parameters.

2.3.3 Testing Hypothesis 1: Moral progress

Table 3. Results of multilevel logistic regressions with log odds of “possibility of progress” as response variable; the morality domain is the reference category

	Abstract		Concrete	
Coefficients	Model I	Model II	Model III	Model IV
Intercept	-.14 (.17)	-.01 (.26)	-.83* (.01)	.13 (.80)
Science	2.13* (.27)	2.16* (.26)	1.01* (.11)	1.01* (.11)
Preferences	.11 (.22)	.12 (.21)	-.51* (.12)	-.51* (.12)
Conventions	.02 (.22)	.02 (.21)	-.24* (.12)	-.24* (.11)
Male		.02 (.22)		-.66 (.80)
Female		reference cat.		-.68 (.81)
Age		-.008 (.01)		-.02 (.01)
Level 2 standard deviation	1.00 (.15)	1.02 (.15)	1.09 (.08)	1.08 (.08)
DIC	1048	1050	3903	3902

Note: Gibbs sampler run with 3 chains, each with 5000 iterations (first 2500 burn-in); Rhat < 1.1 for all parameters; reported coefficients are estimated posterior means; estimated standard errors in brackets; * means estimated 95%-CI does not contain 0; 864 statement evaluations (level 1) nested in 216 participants (level 2) in the abstract models; 3474 statement evaluations (level 1) nested in 193 participants (level 2)

For both the abstract and the concrete versions, we estimate two models: models I and III without control variables, and models II and IV with gender and age as control variables. In all models, morality is the reference category. For the abstract version, Model I strongly refutes the first hypothesis: moral statements are evaluated in a significantly less objectivist manner than are scientific statements ($b = 2.13$, $se = .27$, $p < .05$), and there is no statistical difference in how moral statements are treated compared to statements about personal preferences ($b = .11$, $se = .22$, $p > 0.05$) and statements about social conventions ($b = .02$, $se = .22$, $p > 0.05$). Adding controls for age and gender in Model II, neither of which is significantly associated with the response variable, leaves this pattern unaffected. For the concrete version, Model III

provides a more nuanced picture. Here too, moral statements are evaluated in a significantly less objectivist manner than scientific statements ($b = 1.01$, $se = .11$, $p < .05$). However, contrary to the abstract version, moral statements are evaluated in a more objectivist manner than statements about personal preferences ($b = -.51$, $se = .12$, $p < 0.05$) and statements about social conventions ($b = -.24$, $se = .12$, $p < 0.05$). The inclusion of control variables in Model IV does not change this pattern.

2.3.4 Testing Hypothesis 2: Moral knowledge

Table 4. Results of multilevel logistic regressions with log odds of “possibility of knowledge” as response variable; the morality domain is the reference category

	Abstract		Concrete	
Coefficients	Model I	Model II	Model III	Model IV
Intercept	-.09 (.17)	-.10 (.31)	-.87* (.09)	.34 (.85)
Science	2.14* (.25)	2.13* (.27)	1.80* (.11)	1.80* (.11)
Preferences	.13 (.21)	.11 (.22)	-.58* (.12)	-.58* (.12)
Conventions	.32 (.22)	.30 (.21)	.05 (.11)	.05 (.11)
Male		.06 (.24)		-1.27 (.86)
Female		reference cat.		-1.02 (.86)
Age		-.001 (.01)		-.003 (.01)
Level 2 standard deviation	1.23 (.16)	1.27 (.15)	1.00 (.08)	1.00 (.08)
DIC	1018	1019	3876	3874

Note: Gibbs sampler run with 3 chains, each with 5000 iterations (first 2500 burn-in); Rhat < 1.1 for all parameters; reported coefficients are estimated posterior means; estimated standard errors in brackets; * means estimated 95%-CI does not contain 0; 864 statement evaluations (level 1) nested in 216 participants (level 2) in the abstract models; 3474 statement evaluations (level 1) nested in 193 participants (level 2)

We estimate two models for both the abstract and concrete versions: models I and III without control variables, and models II and IV with gender and age as control variables. In all models, morality is the reference category. For the

abstract version, Model I strongly refutes hypothesis 1: moral statements are evaluated in a less objectivist manner than are scientific statements ($b = 2.14$, $se = .25$, $p < .05$), and there is no difference with personal preferences ($b = .13$, $se = .21$, $p > .05$) or statements about social conventions ($b = .32$, $se = .22$, $p > .05$). Including the controls age and gender in Model II do not significantly change the outcomes. Similar as for the results on moral progress, the results for the concrete version in Model III are slightly more nuanced. Moral statements are evaluated in a significantly less objectivist manner than are scientific statements ($b = 1.80$, $se = .11$, $p < .05$). Contrary to the abstract version, moral statements are evaluated in a more objectivist manner than are statements about personal preferences ($b = -.58$, $se = .12$, $p < .05$). This pattern is similar in Model IV with controls.

2.3.5 Testing Hypothesis 3: Moral error

Table 5. Results of multilevel logistic regressions with log odds of “possibility of error” as response variable; the morality domain is the reference category

	Abstract		Concrete	
	Model I	Model II	Model III	Model IV
Coefficients				
Intercept	-.25 (.18)	-.61* (.29)	-1.15* (.10)	.004 (.85)
Science	1.86* (.26)	1.84* (.22)	.97* (.11)	.98* (.11)
Preferences	-.10 (.22)	-.14 (.22)	-.31* (.13)	-.30* (.12)
Conventions	.33 (.22)	.29 (.21)	.03 (.12)	.03 (.12)
Male		-.22 (.23)		-.87 (.86)
Female		reference cat.		-.95 (.86)
Age		.03* (.01)		-.02 (.01)
Level 2 standard deviation	1.17 (.16)	1.15 (.16)	1.00 (.08)	1.00 (.08)
DIC	1052	1048	3774	3772

Note: Gibbs sampler run with 3 chains, each with 5000 iterations (first 2500 burn-in); Rhat < 1.1 for all parameters; reported coefficients are estimated posterior means; estimated standard errors in brackets; * means estimated 95%-CI does not contain 0; 864 statement evaluations (level 1) nested in 216 participants (level 2) in the abstract models; 3474 statement evaluations (level 1) nested in 193 participants (level 2)

We estimate two models for both the abstract and concrete versions: models I and III without control variables, and models II and IV with gender and age as control variables. In all models, morality is the reference category. For the abstract version, Model I refutes hypothesis 1: moral statements are evaluated in a less objectivist manner than scientific statements ($b = 1.86$, $se = .26$, $p < .05$), and there is no difference with personal preferences ($b = -.10$, $se = .22$, $p > .05$) or statements about social conventions ($b = .33$, $se = .22$, $p > .05$). Including the controls age and gender in Model II do not significantly change the outcomes. Similar as for the results on moral progress and moral error, the results for the concrete version in Model III are more nuanced. Moral statements are evaluated in a significantly less objectivist manner than scientific statements ($b = .97$, $se = .11$, $p < .05$). Contrary to the abstract version, moral statements are evaluated in a more objectivist manner than are statements about personal preferences ($b = -.31$, $se = .13$, $p < .05$). This pattern is similar in Model IV with controls.

2.3.6 Hypothesis tests with subjective categorizations

Table 6. Results of multilevel logistic regressions with log odds of “objective” as response variable; domains are based on participants’ own categorizations of statements from the concrete version; the moral domain is the reference category

	Progress		Knowledge		Error	
	Model I	Model II	Model III	Model IV	Model V	Model VI
Coefs.						
Intercept	-.68* (.11)	.14 (.82)	-.72* (.10)	.30 (.78)	-1.07* (.11)	.21 (.83)
Science	.58* (.12)	.58* (.11)	1.19* (.11)	1.19* (.11)	.63* (.12)	.64* (.12)
Prefs.	-.34* (.11)	-.34* (.11)	-.38* (.11)	-.38* (.11)	-.11 (.12)	-.10 (.12)
Conv.	-.43* (.11)	-.43* (.11)	-.18 (.11)	-.17 (.11)	-.13 (.11)	-.12 (.11)
Male		-.52 (.82)		-1.07 (.75)		-1.02 (.82)
Female		-.56 (.83)		-.86 (.75)		-1.12 (.83)
Age		-.02 (.01)		-.003 (.01)		-.01 (.01)
Level 2 stand-ard deviation	1.06 (.08)	1.05 (.08)	.93 (.08)	.93 (.07)	.99 (.08)	.98 (.08)
DIC	3968	3968	4050	4050	3832	3831

Note: Gibbs sampler run with 3 chains, each with 5000 iterations (first 2500 burn-in); $R_{\text{hat}} < 1.1$ for all parameters; reported coefficients are estimated posterior means; estimated standard errors in brackets; * means estimated 95%-CI does not contain 0; 3474 statement evaluations (level 1) nested in 193 participants (level 2)

We also calculated models for moral progress, knowledge, and error, based on people's own categorizations of statements in the domains of science, morality, personal preferences, and social conventions. Those results are largely similar to the results without subjective categorizations.

3. Discussion

Previous research has investigated whether or not people believe that moral statements are true or false and whether or not moral disagreements imply that only one party can be correct (i.e. as discussed in Chapter 2, and in Goodwin and Darley 2008, Wright et al. 2013, 2014, Pölzler 2018; Pölzler and Wright forthcoming). Those results do not provide evidence in favor of FMO. In light of our discussion in Chapter 3, we should not interpret those findings as providing evidence against FMO either. In this paper, we used a different methodology to investigate whether FMO is true, namely by testing whether people believe in objective moral progress, knowledge, and error.

Our results show that people believe that science (abstractly) is, and scientific statements (concretely) are, a matter of objectivity. For progress, knowledge, and error, people treat morality significantly differently. For all the abstract questions, morality is believed to be on a par with personal preferences and conventions. With regard to the concrete statements, results are slightly more nuanced. For the concrete questions, people seem to believe that moral statements are less objective than scientific statements but more objective than statements about personal preferences or social conventions.

In Chapter 1 and 3, I note that there is a difference between implicit and explicit metaethical commitments and that some philosophers believe

that FMO applies to implicit commitments. If FMO applies to implicit commitments only, the results of this research are only relevant for FMO if our measurement instruments elicit implicit commitments. If our results indeed capture people's implicit commitments, FMO appears to be false. However, some of our measurement instruments are fairly explicit and it is therefore possible that we measure people's explicit commitments. If so and if FMO does not apply to explicit commitments, our results may not provide evidence for FMO but still be interesting for philosophers and psychologists interested in people's explicit commitments. However, future research should investigate whether our measurement instruments elicit implicit or explicit commitments. I therefore conclude that at this point we cannot be sure that this research provides evidence for FMO.

4. Conclusion

In this paper, we investigated whether people believe in the possibility of objective moral progress, knowledge, and error. We assumed that people treat science, abstractly, and scientific statements concretely, as if this domains concerns objective facts. We also assumed that if people believe in the objectivity of moral judgments, the results we obtain for science and morality both abstractly and concretely, would be similar. Our results show that people treat science, both abstractly and concretely, as if objectivism is true for that domain. However, our results also show that they treat morality differently. The results of our study do not support FMO. However, whether they constitute evidence against FMO depends on the nature of FMO, more precisely on whether it concerns implicit or explicit commitments, and how this can be measured.

CONCLUSION

In this PhD thesis, I have investigated whether a specific metaethical assumption is true, namely the thesis of folk moral objectivism (FMO). The aim was to find out whether people believe that moral judgments are objectively true or false. To achieve this, I evaluated whether the results of existing research on folk metaethics support FMO and I conducted empirical studies to investigate whether people are implicitly or explicitly committed to FMO. Overall, this research provides no hard evidence for FMO.

The conclusions of this manuscript are as follows. First, existing research has often used measures that do not directly, or not exclusively, measure moral objectivism, which makes it difficult to assess whether results pertain to FMO. Second, existing research may have measured explicit metaethical commitments while FMO applies to implicit commitments. Third, with regard to the phenomenon of 'perceived objectivity', which is measured by truth-aptness and/or disagreement tasks, it is not clear whether this reflects explicit or implicit commitments. However, given the experimental designs that are employed, which often contain the repetitive use of explicit metaethical language, it is plausible that perceived objectivity reflects explicit commitments. Fourth, empirical studies in Chapter 2 suggest that there are at least three dimensions of thinking about moral truth that underlie perceived objectivity. Moreover, people's scores on those dimensions are differently related to tolerance and a willingness to harm measure. Finally, research conducted as part of this thesis provides some support for the idea that people are implicitly moral objectivists but these results are not conclusive. To obtain a better picture of the results of this research, I will discuss each chapter separately.

In the first chapter, I discussed existing research on folk metaethics. The results suggest that people are not univocally committed to FMO. In fact, there seems to be hardly any participant who consistently treats moral statements as if moral objectivism is true. More specifically, the results show large differences between different individuals with regard to FMO. Moreover, individuals seem to treat statements differently based on moral content. Although these results do not make FMO more plausible, they also do not necessarily

provide evidence against FMO. First, the measurement instruments that are used do not directly measure FMO. Second, FMO may apply to implicit commitments and although there is reason to think that many studies measure explicit commitments, it is not clear what type of commitments is measured.

In the second chapter I investigated whether the psychological construct often used to measure FMO known as 'perceived objectivity', is best measured on a single dimension or is multidimensional. To investigate whether this is the case, I conducted an exploratory factor analysis on a range of ethical statements. The results suggest that there are at least three distinct dimensions that underlie people's metaethical judgments, namely Independent Truth, Universal Truth, and Divine Truth. A confirmatory factor analysis confirms the three-factor structure and reliability tests suggest that each dimension forms reliable subscales. The studies that I used to validate the measurement scale suggest that each of these dimensions is differently related to tolerance and a willingness to harm measure. For instance, it seems that scores on Universal Truth increases people's willingness to harm norm-violators while scores on Independent Truth have the opposite effect.

In the third chapter I investigated whether people are implicitly committed to moral objectivism. Some philosophers believe that FMO applies to implicit commitments. They assume that people are committed to moral objectivity even if they explicitly deny that this is the case. I therefore empirically investigated whether there is evidence to support this claim. To elicit people's implicit commitments, I used three tests developed by David Enoch. Enoch claims that people will respond to those tests as if moral objectivism is true, even if they explicitly reject moral objectivism. The results show that for each of those tests, people seem to support the idea that moral objectivism is true. At the same time, we should be careful to interpret these results as support for FMO. For instance, when we combine results of two of Enoch's tests, we observe that only half of his participants respond in a way that supports FMO and this corresponds to the results of research on more explicit measures.

In the fourth chapter I explored alternative ways of measuring FMO together with co-authors Thomas Pölzler and Jacob Dijkstra. We investigated whether people believe in the possibility of moral progress, knowledge, and error. In our experimental design, we administered both abstract and concrete questions and compared the moral domain with the domains of science, personal preferences, and social conventions. For the abstract questions, people were asked whether they believe in the possibility of objective moral progress, knowledge, and error. For the concrete questions, people were administered concrete cases and they were asked whether or not they believe in the possibility of objective moral progress, knowledge and error. Our results show that people tend to believe that science, both abstractly and concretely, is an objective matter. However, people treat morality differently. That is, for the abstract cases moral matters are treated as on a par with personal preferences and social conventions. For the concrete cases, people believe that moral statements are more objective than statements about personal preferences or social conventions.

In short, the empirical results presented in this PhD thesis do not univocally support FMO. The results of existing research, and those presented in Chapter 2 suggest that FMO is false. However, given that the measurement instruments used may elicit explicit commitments and that FMO may apply to implicit commitments (see Chapter 3), we should not interpret those findings as evidence against FMO. We can draw several lessons from the research presented in this PhD-thesis and I will present them as separate recommendations.

First, philosophers should be precise about what they are committed to when they embrace an empirical thesis like FMO, which includes the question of whether or not FMO necessarily pertains to implicit and/or explicit commitments. FMO is usually defended by observations of ordinary moral discourse and appeals to intuitions. However, appearances can be deceiving. Even if empirical researchers have not always used adequate measurement instruments, it is also true that it is not obvious that different measurement

instruments will lead us to conclude that FMO is true. My second recommendation is to be precise about what implicit commitments are, how they are different from explicit commitments, and how they manifest themselves. Hence, we need to develop a clear conceptual distinction between implicit and explicit commitments. My third recommendation is directed to empirical researchers and results from the former recommendation. If we have a clear conceptual distinction between implicit and explicit commitments, we can then assess whether this constitutes an empirical reality. What we need is a measurement instrument that measures implicit and explicit commitments and the distinction between them. If we obtain those tools and it is clear whether FMO applies to implicit and/or explicit commitments, we can assess whether or not FMO is true. We can then also assess whether it is people's implicit and/or explicit commitments that are multidimensional or are associated with people's tolerance judgments and their willingness to harm norm violators.

Appendix A

HOW NOT TO ARGUE ABOUT IS/ OUGHT INFERENCES IN THE COGNITIVE SCIENCES¹⁸

¹⁸ This chapter is co-authored by Katina Quinzelier and accepted as a paper at *Frontiers in Psychology*.

1. Introduction

When scholars problematize is/ought inferences (IOIs), they sometimes refer to Hume's or Moore's fallacy (e.g., Schneider, 2000; Schroyens, 2009; Elqayam and Evans, 2011). Although inferring "ought" from "is" can be problematic, we argue that, in the context of contemporary IOIs in the cognitive sciences, invoking Hume or Moore might be misguided. This is because Hume's and Moore's arguments concern the validity and soundness of *deductive* inferences while in our view contemporary IOIs in the cognitive sciences are better interpreted as *defeasible* inferences.

In order to avoid misinterpretations, we first clarify key concepts in the debate in section Key Concepts. In section Mind the Gap, we revisit Hume's and Moore's arguments against inferring "ought" from "is," and in section A Debate Shackled, we discuss contemporary IOIs in the cognitive sciences.

2. Key Concepts

Participants in the is/ought debate distinguish between descriptive statements and deontic statements. Descriptive statements describe or predict how the world is. Deontic statements prescribe or proscribe how we should act or reason.

While "is" statements are descriptive statements, "ought" statements can be descriptive as well as deontic. For instance, "the streets ought to be wet because it is raining" is a descriptive statement because it predicts that the streets will be wet. Conversely, "If you do not want to get wet, you ought to carry an umbrella," is a deontic statement because it prescribes what you should do. In this comment, we only discuss "ought" statements as deontic statements. Accordingly, we will not discuss inferences from "is" to descriptive "oughts" (cf. Oaksford and Chater, 2009, 2011), but only inferences from "is" to deontic "oughts" (cf. Oaksford and Sellen, 2000; Stanovich and West, 2000).

We describe an is/ought inference as an attempt to evaluate (i.e., fine-tune, develop, arbitrate between) deontic statements on the basis of descriptive statements. The following is an example of an IOI (1):

Premise: More intelligent people are more likely than less intelligent people to make a guess, instead of reason, when solving the Wason Selection Task.
Conclusion: We ought to make a guess, instead of reason, when solving the Wason Selection Task.

This inference can be interpreted as a deductive argument. As such, the conclusion is true if the inference is valid and sound. A deductive inference is valid if the premises logically entail the conclusions, hence, if it is logically impossible for the premises to be true and the conclusion false. In this inference, it is possible that the premise is true while the conclusion is false. Thus, it is deductively invalid.

Soundness takes the actual truth of the premises (and conclusions) into account: An inference is sound if it is valid *and* all of its premises are true. The inference in this example is not sound because it is invalid. However, were it to be valid, it would still be unsound because the premise is false. More intelligent people are in fact more likely than less intelligent people to reason logically when solving the Wason Selection Task (Stanovich and West, 2000).

An inference can also be interpreted as a defeasible argument. Defeasible inferences have several features, two of which are relevant for our argument (cf. Pollock, 1987, 1992). First, the inference can be correct even if it is not deductively valid. Let us illustrate these features on the basis of the following inference (which is not an is/ought inference) (2):

Premise: X looks red to me.
Conclusion: X is red.

Clearly, the premise does not logically entail the conclusion. However, the inference is defeasibly correct because the premise supports the conclusion—most things that look red to me are, in fact, red.

A second feature of defeasible inferences is that, when the inference is correct, it can still be revised in the light of new information. For instance, if we learn that X is a daisy that is illuminated by red lights, which can make things appear red when they are not, we may suggest the following revised inference (3):

Premise 1: X looks red to me.

Premise 2: X is a daisy that is illuminated by red lights, which can make things appear red when they are not.

Conclusion: X is not red.

While correct defeasible inferences can be revised in the light of new information, valid deductive inferences cannot. If the conclusion follows deductively from a (set of) premise(s), it will still follow deductively no matter how many premises we add. (This is termed the monotonicity of deductive logic.)

All this is relevant for is/ought debates. In section Mind the Gap, we argue that Hume's and Moore's arguments concern the validity and soundness of deductive inferences. In section A Debate Shackled, we explain why IOI's in the cognitive sciences are better interpreted—and evaluated—as defeasible inferences.

3. Mind the Gap

Cognitive scientists often fine-tune, develop or arbitrate between models of how people ought to reason on the basis of theories and data of how people do reason (for a discussion and critique, see Elgayam and Evans, 2011). Critics (e.g., Schneider, 2000; Schroyens, 2009; Elgayam and Evans, 2011) claim that some of these cognitive scientists commit Hume's or Moore's fallacy. However, in line with previous interpretations, we contend that Hume's and Moore's fallacies in the first place preclude deductive inferences that are, respectively, not valid and not sound (cf. Schurz, 1997; Pigden, 2010; Quinteler et al., 2011).

It is useful to introduce a caveat here. Hume and Moore formulated their arguments in the context of ethical "oughts." However, in the cognitive

sciences, their arguments are applied to epistemic "oughts." This is acceptable for standard, logical, interpretations of Hume's fallacy, which seem to hold at least for deontic "oughts" in general (Pigden, 2010, p. 240). In contrast, it is unclear if Moore's fallacy applies to the same extent to non-ethical deontic "oughts." For the sake of argument though, we assume that both fallacies also apply to epistemic "oughts."

Let us now review Hume's fallacy. The standard interpretation of Hume's fallacy states that there are no deductively valid inferences whose premises contain no "oughts" and whose conclusions contain (non-trivial) "oughts" (Schurz, 1997; Pigden, 2010, p. 198–242). For example, the following inference is not deductively valid (3):

Premise: It is the case that human beings apply Bayesian reasoning.

Conclusion: It ought to be the case that human beings apply Bayesian reasoning.

This inference is not deductively valid because it is possible that the conclusions are false while the premises are true. In Hume's words, "ought, or ought not, expresses some new relation or affirmation," which is different than the relation being expressed by "is," or "is not" (1739–1740, Book III, Part I, section Key Concepts). When scholars infer "ought" related conclusions from premises that contain only "isses," they commit Hume's fallacy.

However, Hume also argues that we can add a premise—hereafter termed a bridge principle - that connects "is" and "ought." We can for example suggest the following bridge principle: "if more intelligent people apply reasoning X, we ought to apply reasoning X" (cf. Schneider, 2000, commenting on Stanovich and West, 2000). This principle can then be used as a premise (4):

Premise 1: More intelligent people apply Bayesian reasoning.

Premise 2: If more intelligent people apply Bayesian reasoning, we ought to apply Bayesian reasoning.

Conclusion: We ought to apply Bayesian reasoning.

This inference is now deductively valid: if the premises are true, then the conclusion is also true. Hume's fallacy does not preclude the possibility of finding a plausible bridge principle.

In contrast, Moore's fallacy states that deductive IOI's with bridge principles might be valid, but they are never sound. The reason is that, according to Moore, bridge principles can never be true. Moore's argument is that we should find an analytically true bridge principle, one that spells out what descriptive concepts are in the meaning of the deontic concept (Moore, 1988, §1–15). However, pace Moore, this is impossible because deontic concepts are already simple terms; there is nothing in their meaning than the deontic concept itself. Therefore, there are no true bridge principles. Those who define a deontic concept in descriptive terms and then claim that this definition is analytically true, commit Moore's fallacy (*id.*).

To summarize, we hold that Hume's fallacy states that deductive IOI's are never valid without a bridge principle, while Moore's thesis states that deductive IOI's are never sound because there is no true bridge principle.

4. A Debate Shackled

Invoking Hume's and Moore's fallacy to criticize IOI's in the cognitive sciences can be problematic: If, by making an is/ought inference, authors rarely mean to *deduce* deontic "oughts" from "issues," then their IOI's should not be evaluated on the basis of their deductive validity or soundness. Indeed, we argue that it is more charitable to interpret contemporary IOI's in the cognitive sciences as defeasible inferences: Relevant authors (Oaksford and Sellen, 2000; Stanovich and West, 2000; Douven, 2011) point to descriptive reasons that suggest, rather than logically entail, deontic conclusions. Moreover, these authors aim to make correct inferences that are revisable in the light of new information. Let us take a look at these features of contemporary IOI's in the cognitive sciences.

Stanovich and West (2000) seem to endorse the following inference (5):

Premise: Studies show that more intelligent people are more likely than less intelligent people to reason logically in task A.

Conclusion: We ought to reason logically in task A.

Oaksford and Sellen (2000) remark that the following also holds (6):

Premise: Studies show that high schizotypal people are more likely than low schizotypal people to reason logically in task B.

Conclusion: We ought not to reason logically in task B.

Clearly, these inferences are not deductively valid (*cf.* Schneider, 2000). However, these authors never claimed that their premise deductively entails a deontic conclusion. Instead, both Stanovich and West, 2000, p. 645) and Oaksford and Sellen, 2000, p. 691) speak of descriptive information that *suggests* a certain deontic conclusion. Moreover, these arguments are revisable in the light of new information: What if, for instance, both schizotypy and intelligence are positively correlated with logical reasoning in the same task A? In that case, we have to revise our conclusions that we ought to reason logically in task A. Thus, inferences 5 and 6 are better understood as defeasible inferences and ought to be evaluated accordingly.

Douven (2011) likewise suggests that, in certain cases, descriptive information can be used to inform us about deontic statements. He reasons as follows:

(7) Premise: Human beings update on conditionals by applying rule X.

Conclusion: Human beings ought to update on conditionals by applying rule X.

Again, as a deductive inference, this would be invalid. However, Douven (2011) does not seem to have a deductive inference in mind. In his words, the premise again “suggests” the conclusion, and descriptive information leads to an “outline” of norms or, based on the premise, we can go “some way” in accepting the conclusion (253). This can be understood as a first approximation that can be revised. Moreover, there is no mentioning that descriptive premises logically entail a deontic conclusion.

These examples lead us to conclude that IOI's in the cognitive sciences are better interpreted as defeasible inferences than as deductive inferences. As a consequence, their deductive validity and soundness is not at stake. We therefore suggest that, instead of referring to Hume or Moore, critics of is/ought inferences apply evaluation criteria for defeasible inferences (see e.g., Nute, 1997). This conclusion supplements previous work on the is/ought problem. Schurz (in Pigden, 2010; p. 216), for instance, suggests that defeasible conditional norms might provide plausible bridge principles in ethical is/ought inferences. Other authors suggest that defeasible reasoning can solve problems and paradoxes occurring in monotonic deontic logic (e.g., Nute, 1997). However, previous work usually focused on ethical “oughts” rather than epistemic “oughts.” We therefore hope that this paper spurs research on defeasible reasoning with epistemic “oughts.”

Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Acknowledgments

The research that led to this article was partially supported by the Fund for Scientific Research-Flanders (FWO-V). The authors also thank an anonymous reviewer for helpful suggestions.

SAMENVATTING

We oordelen allemaal wel eens over de morele juistheid of toelaatbaarheid van een bepaalde handeling of stand van zaken. Filosofen noemen zulke oordelen “eerste-orde oordelen”. Veel mensen vinden het bijvoorbeeld moreel juist dat de overheid ingrijpt indien een burger fysiek door een medeburger wordt bedreigd. Veel mensen vinden het moreel onjuist als iemand fraude pleegt en daardoor misbruik maakt van anderen. Sommige mensen vinden abortus, euthanasie, en het homoseksueel moreel toelaatbaar en ze vinden de doodstraf, marteling, liegen, en/of vreemdgaan moreel ontoelaatbaar. Filosofen gebruiken de term “normatieve ethiek” voor filosofisch onderzoek naar zulke eerste-orde oordelen. Dit proefschrift gaat echter over een vraagstuk in de metaethiek. De metaethiek houdt zich niet bezig met normatieve vragen, maar met vragen over de status van onze eerste-orde oordelen. Metaethici onderzoeken bijvoorbeeld of onze morele oordelen als doel hebben de wereld te beschrijven zoals die is, en waar of onwaar kunnen zijn, of dat ze alleen uitspraken geven aan onze emoties of gevoelens.

Dat eerste- en tweede-orde vraagstukken uit elkaar lopen blijkt uit het volgende voorbeeld. Stel dat iemand het eerste-orde oordeel stelt dat abortus moreel ontoelaatbaar is. Die persoon kan tegelijk het tweede-orde oordeel hebben dat de uitspraak “abortus is moreel ontoelaatbaar” niet objectief waar of onwaar is. Misschien vindt die persoon dat het morele oordeel uitspraak geeft aan zijn of haar gevoelens.

De opvatting dat er objectieve morele waarheden zijn noemen filosofen “moreel objectivisme”. De vraag of moreel objectivisme waar is behoort tot de metaethiek. Sommige filosofen denken dat er goede argumenten zijn om te vooronderstellen dat morele uitspraken objectief waar of onwaar zijn terwijl anderen het tegenovergestelde betogen.

Veel filosofen denken, interessant genoeg, dat mensen denken dat morele uitspraken objectief waar of onwaar zijn. Hierbij wordt dus voorondersteld dat wij allemaal geloven dat moreel objectivisme waar is. Anders geformuleerd, deze filosofen denken dat mensen morele uitspraken zoals “abortus is verkeerd”, “de doodstraf is ontoelaatbaar”, of “fraude plegen is fout”, beoor-

delen als waar of onwaar op een vergelijkbare wijze als “3 maal 3 is 9”, “schizofrenie is erfelijk”, of “Zwolle ligt ten noorden van Groningen”, namelijk in de zin dat de waarheid van morele uitspraken onafhankelijk is van ons denken. Sommige filosofen denken dat het feit dat mensen dit geloven ondersteuning geeft aan het idee dat er objectief ware morele oordelen bestaan. Het idee dat mensen geloven dat morele oordelen objectief waar of onwaar zijn omschrijf ik als de these van “Folk Moral Objectivism” (FMO).

In dit proefschrift onderzoek ik of er empirisch bewijs is voor FMO. Het achterliggende idee is dat het vinden van empirische evidente ondersteuning geeft aan filosofische theorieën die FMO als uitgangspunt nemen. Indien empirische ondersteuning volkomen ontbreekt dan kan dit betekenen dat de filosofische theorieën die FMO vooronderstellen moeten worden herzien. De bevindingen uit dit proefschrift kunnen in dit geval bijdragen aan de ontwikkeling van een nieuwe visie op de metaethische overtuigingen van mensen en de filosofische implicaties hiervan voor de metaethiek.

In deze samenvatting bespreek ik eerst mijn algemene bevindingen en vervolgens geef ik per hoofdstuk een korte toelichting op ieder afzonderlijk onderzoek. De centrale onderzoeksvraag van mijn proefschrift, namelijk of FMO waar of onwaar is, beantwoord ik als volgt: Bestaand empirisch onderzoek levert geen empirisch bewijs dat FMO waar of onwaar is.

Een van de redenen hiervoor is dat in bestaand onderzoek vaak meetinstrumenten heeft gebruikt die niet altijd tot valide uitkomsten hebben geleid. De experimentele taken die zijn afgenomen sluiten bijvoorbeeld niet uit dat mensen geloven dat morele oordelen subjectief waar of onwaar zijn of dat alle morele oordelen onwaar zijn.

Een andere reden is dat er een onderscheid is tussen impliciete overtuigingen en expliciete overtuigingen. Laat mij dit illustreren met een voorbeeld. Het is mogelijk om expliciet te ontkennen dat moreel objectivisme onwaar is, zoals sommige filosofen in hun theorieën doen, en tegelijk impliciet te denken dat moreel objectivisme waar is. Misschien denk je expliciet dat alle morele oordelen relatief waar of onwaar zijn maar wanneer je

geconfronteerd wordt met een specifiek voorbeeld, bijvoorbeeld met betrekking tot vrouwenbesnijdenis, abortus, euthanasie, enzovoort, dat je reageert alsof moreel objectivisme waar is. Tegelijk is het niet duidelijk of bestand onderzoek impliciete of expliciete overtuigingen hebben gemeten. Als het klopt dat bestand onderzoek expliciete overtuigingen meet én dat FMO van toepassing is op impliciete overtuigingen dan kunnen we op basis van bestand onderzoek niet vaststellen of FMO waar of onwaar is.

Kortom: Vanwege (1) het feit dat bestaande meetinstrumenten niet altijd op valide wijze FMO hebben gemeten, en (2) dat de meeste onderzoeken mogelijk expliciete overtuigingen hebben gemeten terwijl FMO van toepassing is op impliciete overtuigingen, kom ik tot de conclusie dat bestaande studies geen empirisch bewijs leveren voor of tegen FMO.

In het eerste hoofdstuk heb ik onderzocht wat FMO precies inhoudt en of er empirische evidence voor of tegen FMO is. Ik gebruik daarbij drie criteria om FMO te beschrijven. Het eerste criterium is “Independence”. Independence omvat het idee dat de waarheid of onwaarheid van morele uitspraken onafhankelijk is van de mentale toestanden van individuen of de collectieve overtuigingen van sociale groepen. Een morele uitspraak zoals “abortus is niet toelaatbaar” is dus niet waar of onwaar omdat bepaalde individuen, of bepaalde sociale groepen of culturen, geloven dat het waar of onwaar is. Voor sommige niet-morele uitspraken, zoals bijvoorbeeld “Spinazie is lekker” of “je dient met mes en vork te eten”, geldt dat de waarheid van deze uitspraken afhankelijk is van wat specifieke individuen of sociale groepen vinden of voelen. Als FMO waar is dan zouden mensen moeten geloven dat Independence van toepassing is op morele uitspraken.

Het tweede criterium dat ik gebruik is “Exclusion”. Dit criterium is gebaseerd op het idee dat tegenstrijdige morele uitspraken niet tegelijk waar kunnen zijn. Wanneer een individu gelooft dat liegen moreel toelaatbaar is en een ander individu dat liegen moreel ontoelaatbaar is dan kan maximaal een van beide partijen gelijk hebben. Als FMO waar is dan zouden mensen moeten geloven dat Exclusion van toepassing is op morele uitspraken.

De eerste twee criteria beschouw ik als implicaties van FMO. Of het laatste criterium, “Uniformity”, volgt uit FMO laat ik over aan mijn lezers. Dit criterium betreft het idee dat mensen alle morele uitspraken als objectief waar of onwaar beschouwen en daar geen onderscheid in maken. In andere woorden, mensen passen Independence en Exclusion consistent toe op alle morele uitspraken. Sommige filosofen en psychologen beweren dat Uniformity van toepassing is op morele uitspraken en anderen hebben daar hun twijfels bij.

Mijn evaluatie van bestand onderzoek leidt tot de conclusie dat Independence en Exclusion vaak worden geschonden. Mensen lijken niet consistent te geloven dat morele uitspraken waar of onwaar zijn onafhankelijk van de perspectieven van individuen of sociale groepen. Daarnaast lijken mensen toe te laten dat tegenstrijdige morele uitspraken gelijktijdig waar kunnen zijn. Omdat mensen Independence en Exclusion niet consistent toepassen en onderscheid maken tussen verschillende morele uitspraken concludeer ik dat Uniformity ook geschonden is. Daarom lijkt FMO onwaar te zijn. Omdat echter de meetinstrumenten niet consistent Independence en Exclusion meten en ze waarschijnlijk vooral expliciete overtuigingen meten concludeer ik dat we dit niet als empirisch bewijs voor of tegen FMO kunnen beschouwen.

In het tweede hoofdstuk onderzoek ik of een meetinstrument dat gebruikt wordt door empirische onderzoekers om FMO te meten, namelijk “perceived objectivity”, een specifieke opvatting van morele waarheid meet of verschillende opvattingen van morele waarheid meet. Door gebruik te maken van verschillende vormen van factoranalyse op twintig ethische uitspraken laat ik zien dat er in psychologische zin drie verschillende opvattingen van morele waarheid bestaan. De eerste vorm noem ik “Independent Truth” en dit gaat om het idee dat morele uitspraken waar of onwaar zijn in niet-relatieve zin. De tweede vorm noem ik “Universal Truth” en dit gaat om het idee dat morele uitspraak absoluut waar zijn. De derde vorm noem ik “Divine Truth” en dit gaat om het idee dat de waarheid van morele uitspraken afhankelijk is van het bestaan van bevelen van een goddelijke entiteit. Mijn onderzoek laat verder zien dat de scores van mensen op deze afzonderlijke vormen van waarheid

op verschillende wijze gerelateerd zijn aan de *perceived objectivity* maat. Daarnaast toont mijn onderzoek aan dat mensen die hoog scoren op Independent Truth minder, en mensen die hoog scoren op Universal Truth meer, geneigd zijn om iemand die normen overtreedt te bestraffen.

In het derde hoofdstuk onderzoek ik of mensen in impliciete zin gecommiteerd zijn aan moreel objectivisme. Dat doe ik door drie gedachtenexperimenten te gebruiken als inspiratie voor een survey experiment. Deze gedachtenexperimenten zijn door de filosoof David Enoch ontwikkeld. Enoch (2014) betoogt dat zijn gedachtenexperimenten aantonen dat mensen geloven dat moreel objectivisme waar is *zelfs* als zij dit expliciet ontkennen. Het bestaan van deze gedachtenexperimenten gaven mij de uitgelezen kans om te toetsen of mensen inderdaad impliciet gecommiteerd zijn aan moreel objectivisme. Mijn resultaten tonen aan dat een meerderheid bij ieder gedachtenexperiment reageert op een wijze die aansluit bij het idee dat mensen impliciet gecommiteerd zijn aan moreel objectivisme. In die zin ondersteunen de resultaten van dit specifieke onderzoek FMO. Het is echter ook zo dat wanneer de individuele scores worden gecombineerd, het resultaat een meer genuanceerd beeld verschaft: 50% van de respondenten scoort op een wijze die suggereert dat ze impliciet gecommiteerd zijn aan moreel objectivisme. Een van de nuances die ik hierbij wil aanbrengen is dat dit percentage mogelijk daalt indien er in het onderzoek aanvullende vragen worden gesteld van een gelijksoortige aard. Het is ook mogelijk dat betere tests ervoor zorgen dat dit percentage stijgt.

In het laatste hoofdstuk presenteer ik een onderzoek naar de vraag of mensen geloven in de mogelijkheid van objectieve morele vooruitgang, objectieve morele kennis, en het maken van morele vergissingen in objectieve zin. Dit onderzoek heb ik uitgevoerd samen met Thomas Pölzler en Jacob Dijkstra. In het survey experiment dat we hebben uitgevoerd worden participanten toegewezen aan een abstracte conditie of een concrete conditie. In de abstracte conditie beoordeelen zij voor vier verschillende domeinen (moraliteit, wetenschap, persoonlijke voorkeuren, sociale conventies) of objectieve vooruitgang,

kennis, of het maken van vergissingen, mogelijk is voor dat domein. In de concrete conditie beantwoorden zij een soortgelijke vraag voor verschillende concrete voorbeelden. De resultaten tonen aan dat mensen geneigd zijn om te denken dat wetenschap zowel abstract als concreet een objectieve kwestie betreft. Dit geldt echter niet voor moraliteit. De resultaten van de metingen in de abstracte conditie tonen aan, voor morele vooruitgang, morele kennis, en het maken van morele vergissingen, dat mensen dit niet als objectief beschouwen. De resultaten van de metingen in de concrete conditie tonen ook aan dat wetenschap als objectief wordt beschouwd en moraliteit niet. Tegelijk is het zo dat voor de concrete voorbeelden mensen morele kwesties als iets meer objectief beschouwen dan kwesties over persoonlijke voorkeuren of sociale conventies. Gegeven dat deze resultaten mogelijk inzicht geven in expliciete overtuigingen terwijl FMO van toepassing is op impliciete overtuigingen kunnen we dit niet direct als evidentie tegen FMO beschouwen. Maar dit verschaft ook geen evidentie voor FMO.

De resultaten van mijn proefschrift verschaffen dus geen overtuigend empirisch bewijs voor FMO. In zoverre filosofen geneigd zijn te denken dat mensen expliciet gecommiteerd zijn aan moreel objectivisme lijken zij zich te vergissen. De resultaten in het derde hoofdstuk suggereren dat mensen geneigd zijn in impliciete zin te denken dat moreel objectivisme waar is, maar er is meer onderzoek naar impliciete overtuigingen nodig voordat we dit stellig kunnen concluderen. In toekomstig onderzoek zou onderzocht kunnen worden hoe impliciete en expliciete overtuigingen verschillen. Verder kan er een meetinstrument ontwikkeld worden dat beide overtuigingen meet. Op basis daarvan kan onderzocht worden of mensen impliciet en/of expliciete gecommiteerd zijn aan moreel objectivisme en of FMO waar is. Ook kan dan worden onderzocht of empirische bevindingen over tolerantie, en de bereidheid om anderen te straffen als zij een norm overschrijden, in relatie staat tot impliciete overtuigingen, tot expliciete overtuigingen, of beide.

CURRICULUM VITAE

Lieuwe Zijlstra was born in Hoogezand-Sappemeer, the Netherlands, on April 28, 1983. He studied Sociology and Philosophy at the University of Groningen and received his Research Master degrees in 2013 (cum laude). At the end of his studies, he was awarded a Huygens Scholarship to study at Yale University in New Haven (CT) in the United States. After his studies, he enrolled in a PhD program at Ghent University with Prof. dr. Jan Verplaetse as his promotor. The first year of his PhD was spent on researching the topic of folk metaphysics and, more generally, the relationship between empirical research and philosophical theorizing at Ghent University. In this year, he also co-authored the chapter in Appendix A together with Dr. Katinka Quintelier, which was published in *Frontiers in Psychology*. The second year was spent as a Visiting Assistant in Research at the Faculty of Philosophy and the Department of Psychology of Yale University with Prof. Joshua Knobe as his advisor. At Yale, he developed and tested the measurement scale that is presented in chapter 2 of this PhD thesis and the studies needed to validate the scale, which is described in a journal article that is published in *Journal of Experimental Social Psychology*. The final two years of his PhD were spent at the Faculty of Philosophy of the University of Groningen with Prof. dr. Frank Hindriks as promotor and Dr. Daan Evers as co-promotor. In these years, he developed the research on implicit commitments (chapter 3) and the possibility of objective moral progress, knowledge and error together with Dr. Thomas Pölzler and Dr. Jacob Dijkstra (chapter 4). He is currently employed as a lecturer “Social Sciences” in the liberal arts program of the University College Groningen.

BIBLIOGRAPHY

- Ayer, A.J. (1936). *Language, Truth and Logic*. Oxford, UK: Ryerson Press
- Bartels, D. M., Bauman, C.W., Cushman, F.A., Pizarro, D.A., & McGraw, A.P., (2015). Moral judgment and Decision Making. In G. Keren & G. Wu (Eds.) *The Wiley Blackwell Handbook of Judgment and Decision Making*. Chichester, UK: Wiley.
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67 (1), 1-48. doi: 10.18637/jss.v067.i01
- Beebe, J. R., Qiaoan, R., Wysocki, T., & Endara, M.A. (2015) Moral Objectivism in Cross-Cultural perspective. *Journal of Cognition and Culture*, 15, 386-401.
- Beebe, J. R. and Sackris, D. (2016). Moral Objectivism Across the Lifespan. *Philosophical Psychology*, 29 (6), 912-929.
- Bentler, P.M. (1990). Comparative Fit Indexes in Structural Models. *Psychological Bulletin*, 107 (2), 238-246.
- Björnsson, G. (2012). Do 'Objectivist' Features of Moral Discourse and Thinking Support Moral Objectivism? *Journal of Ethics* 16 (4): 367-93. doi: 10.1007/s10892-012-91319
- Blackburn, S. (1993). *Essays in Quasi-Realism*. New York: Oxford University Press.
- Brink, D. (1984). Moral Realism and the Sceptical Arguments from Disagreement and Queerness. *Australasian Journal of Philosophy*, 62, 111-125.
- Brink, D. (1989). *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Boyd, R. (1988). *Essays on Moral Realism*. In G. Sayre-McCord (ed.) *Essays on Moral Realism*. Ithaca: Cornell University Press.
- Browne, M.W. & Cudeck, R. (1993). Alternative Ways of Assessing Model Fit. In K.A. Bollen & J.S. Long (Eds.), *Testing Structural Equation Models* (pp. 136-162). Beverly Hills, CA: Sage.
- Buhrmester, M., Kwang, T., & Gosling, S.D. (2011). Amazon's Mechanical Turk: A New Source of Inexpensive, yet High-Quality Data? *Perspectives on Psychological Science*, 6(1), 3-5.
- Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences* (2nd ed.). Hillsdale, NJ: Lawrence Earlbaum Associates.
- Cuneo, T. (2007). *The Normative Web: An Argument for Moral Realism*. Oxford: Oxford University Press.
- Dancy, J. (1986). Two Conceptions of Moral Realism. *Proceedings of the Aristotelian Society*, Supplemental Volume 60, 167-187
- Douven, I. (2011). A Role for Normativism. *Behavioral and Brain Sciences*, 34, 252-253. doi: 10.1017/S0140525X11000471
- Dreier, J. (1990). Internalism and Speaker Relativism. *Ethics*, 101(1), 6-26
- Elqayam, S., and Evans, J. St. B. T. (2011). Subtracting "ought" from "is": de scriptivism versus norma- tivism in the study of human thinking. *Behavioral Brain Sciences*, 34, 233-248. doi: 10.1017/S0140525X1100001X
- Enoch, D. (2005). Why Idealize? *Ethics*, 115, 759-787
- Enoch, D. (2014). Why I am an Objectivist About Ethics (and Why You Are, Too). In: R. Shafer Landau, ed. *The Ethical Life: Fundamental Readings in Ethics and Moral Problems*, 192-205. New York: Oxford University Press
- Fabrigar, L.R., Wegener, D.T., MacCallum, R. C., & Strahan, E. J. (1999). Evaluating the Use of Exploratory Factor Analysis in Psychological Research. *Psychological Methods*, 4 (3), 272-299. doi: 10.1037/1082-989X.4.3.272
- Faul, F., Erdfelder, E., Buchner, A., & Lang, A. G. (2009). Statistical Power Analysis using G*Power 3.1: Tests for Correlation and Regression Analyses. *Behavior Research Methods*, 41, 1149-1160. doi: 10.3758/brm.41.4.1149
- Feltz, A. & Cokely E.T. (2008). The Fragmented Folk: More Evidence of Stable Individual Differences in Moral Judgment and Folk Intuitions. In B.C. Love, K. McRae and V.M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 1771-1776). Austin, TX: Cognitive Science Society
- Fisher, M., Knobe, J., Strickland, B. & Keil, F.C. (2017) The Influence of Social Interaction on Intuitions of Objectivity and Subjectivity. *Cognitive Science*, 41 (4), 1119-1134

- Forsyth, D. R. (1980). A Taxonomy of Ethical Ideologies. *Journal of Personality and Social Psychology*, 39, 175-184
- Goodwin, G.P., & Darley, J.M. (2008). The Psychology of Metaethics: Exploring Objectivism. *Cognition*, 106, 1339-1366. doi: 10.1016/j.cognition.2007.06.007
- Goodwin, G. P. and Darley, J. M. 2010: The Perceived Objectivity of Ethical Beliefs:
- Psychological Findings and Implications for Public policy. *Review of Philosophy and Psychology*, 1, 1-28.
- Goodwin, G.P., & Darley, J. (2012). Why Are Some Moral Beliefs Seen as More Objective Than Others? *Journal of Experimental Social Psychology*, 48, 250-256. doi: 10.1016/j.jesp.2011.08.006
- Graham, J., & Haidt, J. (2010). Beyond Beliefs: Religions Bind Individuals into Moral Communities. *Personality and social psychology review*, 14(1), 140-150.
- Graham, J., Meindl, P., Beall, E., Johnson, K. M., & Zhang, L. (2016). Cultural Differences in Moral judgment and Behavior, Across and Within societies. *Current Opinion in Psychology*, 8, 125-130.
- Haidt, J., Koller, S., & Dias, M. (1993). Affect, Culture, and Morality, or is it Wrong to Eat Your Dog? *Journal of Personality and Social Psychology*, 65, 613-628.
- Haidt, J. Rosenberg, E., & Hom, H. (2003). Differentiating diversities: Moral Diversity is Not Like Other Kinds. *Journal of Applied Social Psychology*, 33, 1-36
- Hare, R.M. (1954). Universalizability, *Proceedings of Aristotelian Society*, 55, 295 -312
- Harmann, G. (1975). Moral Relativism Defended. *The Philosophical Review*, 84, 3-22
- Harmann, G. (2015). Moral Relativism is Moral Realism. *Philosophical Studies*, 172, 855-863

- Hu, L. & Bentler, P.M. (1999). Cutoff Criteria for Fit Indexes in Covariance Structure Analysis: Conventional Criteria Versus New Alternatives. *Structural Equation Modeling: A Multidisciplinary Journal*, 6 (1), pp. 1-55. <http://doi.org/10.1080/10705519909540118>
- Huemer, M. (2005). *Ethical Intuitionism*. Basingstoke: Palgrave MacMillan
- Hume, D. (1739-1740). *A Treatise of Human Nature*. Available online at: <http://www.gutenberg.org/ebooks/4705>
- Joyce, R. (2006). *The Evolution of Morality*. Cambridge, MA: MIT Press
- Kant, I. (1959). *Foundations of the Metaphysics of Morals* (L.W. Back, Trans.) Indianapolis, IN: Bobbs-Merrill (Original work published 1789)
- Kaiser, H.F. (1974). An Index of Factorial Simplicity. *Psychometrika*, 39, 31-36.
- Mackie, J. (1977). *Ethics: Inventing Right and Wrong*. New York: Penguin
- McNaughton, D. (1988). *Moral Vision: An Introduction to Ethics*. Oxford: Blackwell
- Moore, G. E. (1988). *Principia Ethica*. Available online at: <http://fair-use.org/g-e-moore/principia-ethica/> (Accessed: March, 2014).
- Murphy, M.C. (1998). Divine Command, Divine Will and Moral Obligation. *Faith and Philosophy*, 16 (1)
- Murray, D. & Nahmias, E. (2014). Explaining Away Incompatibilist Intuitions. *Philosophy and Phenomenological Research*, 88(2): 434-467. doi: 10.1111/j.1933-1592.2012.00609.x
- Netemeyer, R.G., Bearden, W.O., & Sharma, S. (2003). *Scaling Procedures. Issues and Applications*. Thousand Oaks: Sage Publications.
- Nichols, S. (2004). After Objectivity: An Empirical Study of Moral Judgment. *Philosophical Psychology*, 17 (1), 3-26. doi: 10.1080/0951508042000202354
- Nichols, S. & Knobe, J. (2007). Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions. *Nous* 41(4): 663-685
- Nucci, L., Turiel, E., & Enarnacion-Gawrych, G. (1983). Children's Social Interactions and Social Concepts in the Virgin Islands. *Journal of Cross-Cultural Psychology*, 14, 469-487
- Nute, D. (1997). *Defeasible Deontic Logic*. New York: Springer Verlag

- Oaksford, M., and Chater, N. (2011). The “is-ought fallacy” fallacy. *Behavioral Brain Sciences*, 34, 262–263. doi: 10.1017/S0140525X11000665
- Oaksford, M., and Sellén, J. (2000). Paradoxical individual differences in conditional inference. *Behavioral Brain Sciences*, 23, 691–692.
- Paolacci, G., & Chandler, J. (2014). Inside the Turk: Understanding Mechanical Turk as a Participant Pool. *Current Directions in Psychological Science*, 23, 184–188. doi: 10.1177/0963721414531598
- Piazza, J., & Landy, J.F. (2013). “Lean Not on Your Own Understanding”: Belief That Morality is Founded on Divine Authority and Non-Utilitarian Moral Judgments. *Judgment and Decision-Making*, 8(6), 639–661
- Pigden, C. R. (2010). *Hume on Is and Ought*. New York: Palgrave Macmillan.
- Pollock, J.L. (1987). Defeasible Reasoning. *Cognitive Science*, 11(4), 481–518
- Pollock, J. L. (1992). How to reason defeasibly. *Artificial Intelligence*, 57, 1–42. doi: 10.1016/0004-3702(92)90103-5
- Pölzler, T. (2018). How to Measure Moral Realism. *Review of Philosophy and Psychology*, 9 (3): 647–670
- Quinn, P. L. (1978). *Divine Commands and Moral Requirements*. Oxford: Clarendon Press.
- Quintelier, K. J. P., De Smet, D., & Fessler, D. M. (2013). The Moral Universalism-Relativism Debate. *Klasis Revue Philosophique*, 27, 211–262.
- Quintelier, K.J.P., Van Speybroeck, L., & Braeckman, J. (2011). Normative Ethics Does Not Need a Foundation: It Needs More Science. *Acta Biotheoretica*, 59, 29–51. doi: 10.1007/s10441-010-9096-7.
- Quintelier K.J.P., & Zijlstra, L. (2014) How Not to Argue About Is/Ought Inferences in the Cognitive Sciences. *Frontiers in Psychology*. doi: 10.3389/fpsyg.2014.00503.
- Rai, T., & Holyoak, S. (2013). Exposure to Moral Relativism Compromises Moral Behavior. *Journal of Experimental Social Psychology*, 49, 995–1001. doi: 10.1016/j.jesp.2013.06.008
- Raitlon, P. (1986). Moral Realism. *Philosophical review*, 95(2), 163–207
- Rosseel, U. (2012) Lavaan: an R Package for Structural Equation Modeling. *Journal of Statistical Software*, 48 (1), 1–36.
- Sarkissian, H., Park, J., Tien, D., Wright, J.C., & Knobe, J. (2011). Folk Moral Relativism. *Mind & Language*, 26, 482–505. doi: 10.1111/j.1468-0017.2011.01428.x.
- Sarkissian, H., & Phelan, M. (2019). Moral Objectivism and a Punishing God. *Journal of Experimental Social Psychology*, 80, 1–7
- Saroglou, V. (2011). Believing, Bonding, Behaving, and Belonging: The Big Four Religious Dimensions and Cultural Variation. *Journal of Cross-Cultural Psychology*, 42(8), 1320–1340.
- Shafer-Landau, R. (2003). *Moral Realism: A Defence*. Oxford: Oxford University Press
- Schneider, S. L. (2000). An Elitist Naturalistic Fallacy and the Automatic-Controlled Continuum. *Behavioral Brain Sciences*, 23, 695–696. doi: 10.1017/S0140525X00 553436.
- Schroyens, W. (2009). On Is And Ought: Levels of Analysis and the Descriptive Versus Normative Analysis of Human Reasoning. *Behavioral Brain Sciences*, 32, 101–102. doi: 10.1017/S0140525X09000478.
- Shweder, R.A. (2002). ‘What About Femal Genital Mutilation?’ and Why Culture Matters in the First Place. In R. Shweder, M. Minow, & H. Markus (Eds.) *Engaging cultural differences: the multicultural challenge in liberal democracies*.
- Schneider, S. L. (2000). An elitist naturalistic fallacy and the automatic-controlled continuum. *Behavioral Brain Sciences*, 23, 695–696. doi: 10.1017/S0140525X00 553436
- Schurz, G. (1997). *The Is-Ought Problem—An Investigation in Philosophical Logic*. Dordrecht: Kluwer.
- Skitka, L. J., & Mullen, E. (2002). Understanding Judgments of Fairness in a Real-World Political Context: A Test of the Value Protection Model of Justice Reasoning. *Personality and Social Psychology Bulletin*, 28, 1419–1429
- Skitka, L. J., Bauman, C. W., & Sargis, E. G. (2005). Moral Conviction: Another Contributor to Attitude Strength or Something More? *Journal of Personality and Social Psychology*, 88, 895–917

- Smetana, J. G. (1981). Preschool Children's Conceptions of Moral and Social Rules. *Child Development*, 52(4), 1333-1336. doi: 10.2307/1129527
- Smetana, J. G. (1983). Social-Cognitive Development: Domain Distinctions and Coordinations. *Developmental Review*, 3(2), 131-147. doi: 10.1016/0273-2297(83)90027-8
- Smetana, J. G., & Braeges, J. L. (1990). The Development of Toddler's Moral and Conventional Judgments. *Merrill-Palmer Quarterly*, 36(3), 329-346
- Smith, M. (1994). *The Moral Problem*. Oxford: Blackwell
- Stanovich, K. E., and West, R. F. (2000). Advancing the rationality debate. *Behavioral Brain Sciences*, 23, 701-717. doi: 10.1017/S0140525X00623439
- Stevenson, C. L. (1944). *Ethics and language*. New Haven, CT: Yale University Press.
- Stevenson, C.L. (1963). *Facts and Values*, New Haven, CT: Yale University Press
- Streumer, B. (2017). *Unbelievable Errors: An Error Theory About All Normative Judgments*. Oxford: Oxford University Press doi: 10.1937/oso/9780198785897.001.0001
- Sturgeon, N. (1986). Moral Explanations. In G. Sayre-McCord (ed.) *Essays on Moral Realism*. Ithaca: Cornell University Press
- Swain, S., Weather, D., & Niedrich, R. (2008). Assessing Three Sources of Misresponse to Reversed Likert Items. *Journal of Marketing Research*, 45, 116-131
- Tabachnick, B. G., & Fidell, L. S. (2001). *Using Multivariate Statistics*. Boston: Allyn and Bacon.
- The United Nations (1948). *Universal Declaration of Human Rights*
- Turiel, E. (1983). *The Development of Social Knowledge: Morality and Convention*. New York, NY: Cambridge University Press
- Turiel, E. (1998). The Development of Morality. In W. Damon & N. Eisenberg (Eds.), *Handbook of child psychology: Social, Emotional, and Personality Development* (pp. 863-932). Hoboken, NJ: John Wiley & Sons Inc.
- Velicer, W.F. & Fava, J.L. (1998). Effects of Variable and Subject Sampling on Factor Pattern Recovery. *Psychological Methods*, 3, 231-251

- Wong, D.B. (2006). *Natural Moralities: A Defense of Pluralistic Relativism*, New York: Oxford University Press.
- Wright, J.C., Cullum, J., & Schwab, N. (2008). The Cognitive and Affective Dimensions of Moral Conviction: Implications for Attitudinal and Behavioral Measures of Interpersonal Tolerance. *Personality and Social Psychology Bulletin*, 34 (11), 1461-76. doi: 10.1177/0146167208322557
- Wright, J.C., McWhite, C., & Grandjean, P. (2014). The Cognitive Mechanisms of Intolerance: Do our Metamoral Commitments Matter? In T.Lombrozo, J. Knobe, and S. Nichols (Eds.), *Oxford Studies in Experimental Philosophy, Volume 1* (pp. 28-61). Oxford University Press
- Yilmaz, O. & Bahçekapılı, H. G. (2015). Without God Everything is Permitted? The Reciprocal Influence of Religious and Meta-Ethical Beliefs. *Journal of Experimental Social Psychology*, 58, 95-100
- Young, L., & Durwin, A.J. (2013). Moral realism as moral motivation: The Impact of Metaethics on Everyday Decision-Making. *Journal of Experimental Social Psychology*, 49 (2), 302-306. doi: 10.1016/j.jesp.2012.11.013
- Zijlstra, L. (2019). Folk Moral Objectivism and its Measurement. *Journal of Experimental Social Psychology*, 84. <https://doi.org/10.1016/j.jesp.2019.04.005>
- Zijlstra, L. (unpublished). Does Empirical Research Debunk Folk Moral Objectivism? Under review at an international peer-reviewed journal.
- Zijlstra, L. (unpublished). Are People Implicitly Moral Objectivists? Under review at an international peer-reviewed journal.
- Zijlstra, L., Fölzler, T., & Dijkstra, J. (unpublished). Do People Believe in Objective Moral Progress, Knowledge, and Error?