


SCHRÖDINGER'S CATEGORIES:
THE INDETERMINACY OF
FOLK METAETHICS



A Dissertation

*Presented to the Faculty of the Graduate School
of Cornell University
in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy*

by

Lance Spencer Bush
May 2023

© 2023 Lance Spencer Bush

SCHRÖDINGER'S CATEGORIES: THE INDETERMINACY OF FOLK METAETHICS

Lance Spencer Bush, Ph.D.

Cornell University 2023

Metaethics is a field of philosophy that addresses fundamental questions about the nature of morality. One of the central disputes in metaethics is whether *moral realism* is true. Moral realism is the claim that there are *stance-independent moral facts*, moral facts that are true independent of the standards or values of individuals or groups, much like scientific facts (e.g., the shape of the earth) aren't made true by personal preference or cultural consensus. *Moral antirealism* is the claim that *there are no stance-independent moral facts*. Research on *folk metaethics* studies whether ordinary people (i.e., *nonphilosophers*) endorse realism or antirealism, or speak and think in ways that commit them to one of these views. Some researchers maintain that nearly everyone endorses either realism or antirealism, but not both. Yet most research suggests significant interpersonal and intrapersonal variation in *folk metaethics*: some people are more inclined towards realism, and others antirealism, while most people are *metaethical pluralists*: they are moral realists about some moral issues and antirealists about others. Regardless of the account in question, *all* existing research presumes that there is a *determinate* fact about whether people are realists or antirealists. I argue that existing evidence does not support this conclusion. Instead, the best account of folk metaethics may be *metaethical indeterminacy*: ordinary people are neither realists nor antirealists, and neither best explains the way people speak or think. The case for *metaethical indeterminacy* proceeds in two steps. First, I argue that *all published studies on folk metaethics rely on invalid measures*. Second, I present empirical evidence that challenges the validity of existing research on folk metaethics and supports metaethical indeterminacy. I evaluate the proportion of people who interpret questions about metaethics as intended, using open response questions, as well as multiple choice questions and Likert scale items. These studies show that most people do not interpret questions

about metaethics as researchers intend. I conclude with a study that demonstrates how forced choice paradigms can create the misleading appearance of a genuine pattern of determinate folk philosophical views, even where none plausibly exist.

PREVIEW

BIOGRAPHICAL SKETCH

I have been told that I tend to focus on the negative, and I think that's mostly correct. I'd like to think that this gives me an edge. But it's just who I am. It's not in my nature to provide a positive biographical sketch; the kind that presents an array of highlights and memories like polished trophies. We all know these glossy histories are at best only half-truths. Sure, I've had bright moments. My mom tucking me into bed, folding half the blanket over me and declaring me a "taco." Foggy memories plodding through the swamps of New Jersey with my brother, casting a wary gaze over my shoulder for signs of the Jersey Devil. Winning first place in the pinewood derby. Tearing open a pack of cards from the *Weatherlight* expansion of "Magic: The Gathering" to find a Thundermare staring back at me, eyes ablaze.

Childhood was a time of wonder. But it was also a time of pain, confusion, and loneliness. I'd like to think the person that I became benefited from the hardships, but I fear I'm simply all that remains after so many enervating experiences left me pitted and gouged. I dropped out of high school in tenth grade—a rocky start, but I turned things around. Community college at sixteen. Finished two bachelor's degrees, then completed an MA at Tufts. Hobnobbed with students at Harvard and Oxford. Spent a few months as a research assistant in the Morality Lab at Boston College, and a brief stint teaching philosophy before I managed to get into Cornell. I've had a lot of opportunities, to be sure, and no small amount of luck.

But like I said, I tend to focus on the negative. In August 2019 I developed a chronic medical condition that nearly scuttled any hope I had of finishing my dissertation. A few months later, my mom died after a long battle with cancer—she beat the cancer, but she died anyway. I see pictures of myself from 2018 and I don't recognize the smile there anymore, or the absence of gray hair.

This dissertation is the culmination of over ten years of thinking about metaethics, moral

psychology, and moral philosophy. My aims are largely critical: critical of every study on the psychology of metaethics, of moral psychology, and of philosophy in general. If my central arguments are correct, then an entire literature is fundamentally misguided and thoroughly mistaken. Yet I think it'd be a mistake to see my ambitions as merely destructive, or to think that nothing good can come from an emphasis on criticizing the work of others. When we clear away the mistakes and confusions of the past, wherever possible, we make way for something new, something better. But the pursuit of something better is a task for someone else. Because, after all, I tend to focus on the negative.

PREVIEW

To mom

ACKNOWLEDGEMENTS

No self-respecting person would begin their dissertation without an obligatory tipping of the hat to everyone who helped them along the way. As cynical as I am, this is one ritual I wholeheartedly endorse. Since I have many people to acknowledge, I'll begin by listing everyone I mention: Dr. David Pizarro, Dr. Laura Niemi, Dr. Shaun Nichols, Dr. Morten Christiansen, David Moss, Andres Montealegre, Tyler Millhouse, Don Loeb, Michael Gill, Sam Rosen, Brian Tebbitt, Bryce Gessell, Todd K. Shackelford, Sebastian Deri, Rebekah Davis, Thomas Pölzler, Dillon Bowen, Lieuwe Zijlstra, Geoffrey Goodwin, Jennifer Cole Wright, Ross Colebrook, Adam Gibbons, Tyler John, David Newman, Geoffrey Goodwin, and my wife Priscila Serpa.

First, I owe the deepest debt of gratitude to my advisor and the chair of my committee, Dr. David Pizarro. I am grateful for the chance to demonstrate what someone with a fundamentally different approach to psychology can bring to the field and can only hope the gamble pays off. I am also grateful for the approach David took as an advisor and a mentor. I'm an inveterate loner. Stubborn and independent, many of my greatest shortcomings stem from my unwillingness to conform with conventions even when they are there for good reason. David provided the breathing room for me to retain a sense of autonomy, Yet David also knew how and when to step in as it was necessary. There were many times I was afraid to step out of my comfort zone and address the weaknesses in my knowledge and skills. More than merely serving as a mentor, David excelled at creating an atmosphere that acknowledged the pressures of graduate study, but with an ever-present humor that made the whole process tolerable.

I also want to acknowledge the rest of their committee, Dr. Laura Niemi, Dr. Shaun Nichols, Dr. Morten Christiansen, for their insights, thoughts, and contributions. I haven't had the chance to

interact with the rest of the committee as much as I'd have liked, but I am nevertheless grateful for their insights, feedback, and criticisms.

I owe a special thanks to David Moss. David has contributed more than anyone to my research and my development as a thinker. David has been a constant source of insight, moderating my brasher impulses with a more measured approach to research and presenting countless insights and observations that have furthered my knowledge of psychological research and folk philosophy. In fact, although I've refined and implemented the coding scheme employed in chapter 4, it was David Moss who initially proposed this approach. David has also contributed a great deal to the studies reported here, assisting in the formulation of the hypotheses and the measures used, as well as contributing to the many criticisms I level against existing research on folk metaethics. David is in every way a research collaborator and deserves much of the credit for the final product.

I also want to thank my other colleagues and collaborators. Andres Montealegre has been a friend, research colleague, and supporter of the direction of my research for many years now. We have worked on numerous projects, and I have learned more from Andres about how to use R, quality research design, open science practices, the importance of a rigorous approach to survey design, preregistration, and much more than I ever could have imagined.

Thanks, as well, to Tyler Millhouse. Tyler was my roommate when I was a philosophy graduate student at Tufts. Our many discussions about philosophy both deepened my understanding of metaethics and broadened the scope of my knowledge with other areas of philosophy. Tyler also tempered many of my worst philosophical habits and served as a role model for the ability to approach philosophical disagreements with calm, composure, and candor. Tyler has provided feedback on several projects related to folk metaethics and has worked with me on many of my earliest projects, and he continues on occasion to work with me on current projects.

Don Loeb has also earned my gratitude. Don's work has had a profound influence on me, inspiring and motivating me to continue my pursuits with the confidence that a graduate student without the academic credentials of professional philosophers wasn't completely misguided in thinking there was something wrong with the methods and conclusions of contemporary metaethics. Don has also been an open book about his life, sharing many stories and anecdotes. His avuncular demeanor and his ability to befriend virtually everyone he meets is a reminder that philosophers need not be insular and aloof. I have yet to learn how to be as kind and compassionate towards others, and I often fear that it's beyond me. Thanks for being so awesome, Don!

I also want to thank Michael Gill. Dr. Gill's paper on indeterminacy and variability in metaethics could be seen as the key paper motivating my research. This is the only paper I know of that explicitly argued for metaethical indeterminacy. In many ways, this project is a very long footnote to Loeb's incoherentism and Gill's Indeterminacy-variability thesis.

I want to thank Sam Rosen for suggesting the title for the paper. I love bad jokes and terrible wordplay, and *Schrödinger's Categories* scratched that itch perfectly.

Thanks also to Brian Tebbitt, who provided valuable feedback on the content and direction of the dissertation and helped remind me to focus rather than go on a bunch of tangents. Thanks also for the many helpful discussions which helped shape the way I think by refining my views about language and the relationship between language, philosophy, and psychology.

There are far too many other people to thank as well. Thanks to the editors and reviewers at *New Ideas in Psychology* and *Diametros*. A special thanks to Bryce Gessell for helpful discussions and feedback. Thanks to Todd K. Shackelford for inspiring me to pursue psychology in addition to philosophy, for providing so many opportunities to publish articles and to interact with incredible researchers, and most of all for helping me to develop a richer understanding of evolution. If his classes hadn't been so engaging, I may have never moved into my current line of research. Thanks to

Sebastian for extensive help with R and feedback on how to write a dissertation. Thanks to everyone who attended the folk moral relativism workshop and to Thomas Pözlner for organizing it. Thank you to Rebekah Davis for providing critical feedback on the dissertation and for encouragement with my research. Thanks to many others for their helpful discussions and feedback, including Dillon Bowen, Lieuwe Zijlstra, Geoffrey Goodwin, Jennifer Cole Wright, Ross Colebrook, Andrew Maul, Hagop Sarkissian, Adam Gibbons, Tyler John, David Newman. I'd also like to thank Geoff Goodwin for generously sharing data from his research, despite my critical analysis of portions of it. Thank you to everyone else at Cornell who provided support or feedback or simply attended one of my talks. Thanks also to the many people I don't remember or never met who saw one of my posters or presentations at conferences. Thanks also to my parents, who supported me the whole way. Thanks to our wonderful dog Tchutchu, who is sadly no longer with us. She was my sole company for most of the writing process and throughout the solitude of the pandemic. I'd like to think she did the most to motivate me.

Finally, I'd like to thank my wife, Priscila. She has had to put up with all of my flaws and shortcomings the most, yet she's done more than everyone else to keep my head in the game and my fingers typing. Priscila undoubtedly deserves the most credit for seeing this work to completion. Priscila helped with formatting my figures and tables and assisted with numerous aspects of the text, from feedback on the overall organizational structure, deciding which studies to include, providing an outsider's perspective on studies and how I frame my arguments, to virtually every element of writing, formatting, and thinking about the topics I discuss. Yet all the academic feedback and support in the world would be irrelevant if I didn't have the motivation and focus to finish. A lot of terrible things happened in the past few years that held me back and made finishing this project an immense struggle. Priscila has been with me every step of the way, providing encouragement and support, and pushing

me to get organized, to get focused, and to take care of my mental and physical health. I couldn't have done this without her.

PREVIEW

TABLE OF CONTENTS

BIOGRAPHICAL SKETCH.....	14
ACKNOWLEDGEMENTS.....	17
CHAPTER 1: INTRODUCTION	24
CHAPTER 2: CRITIQUE OF THE DISAGREEMENT PARADIGM	37
2.0 INTRODUCTION	37
2.1 THE DISAGREEMENT PARADIGM	38
2.2 OVERVIEW OF CRITIQUE OF THE DISAGREEMENT PARADIGM.....	42
2.3 INADEQUATE RESPONSE OPTIONS.....	46
2.4 CONFLATIONS WITH UNINTENDED CONCEPTS.....	50
2.5 A SUMMARY OF REMAINING PROBLEMS	62
2.6 CONCLUSION.....	66
CHAPTER 3: CRITIQUE OF ALTERNATIVE PARADIGMS.....	69
3.0 INTRODUCTION	69
3.1 SCALES.....	70
3.2 PROBLEMS OF VALIDITY IN METAETHICS SCALES: AN EXAMPLE.....	71
3.3 TRAINING PARADIGMS	82
3.4 CONCLUSION.....	91
CHAPTER 4: OPEN RESPONSE PARADIGMS.....	92
4.0 INTRODUCTION	92
4.1 RATIONALE FOR USING OPEN RESPONSE DATA.....	93
4.2 GENERAL METHODS AND PROCEDURES.....	99
4.3 GENERAL PREDICTIONS	104
4.4 SUMMARY OF STUDIES	106
4.5 STUDIES 1A & 1B: SOURCE OF DISAGREEMENT IN GOODWIN & DARLEY (2008).....	106
4.6 STUDY 1C.....	117
4.7 STUDY 2: INTERPRETING OTHER PEOPLE’S RESPONSES TO METAETHICS STIMULI.....	117
4.8 STUDY 3: CHARITY & NONCOGNITIVISM.....	127
4.9 STUDY 4: DIRECT REALISM AND RELATIVISM.....	132
4.10 STUDY 5: INTERPRETATION OF THE MORAL RELATIVISM SCALE (MRS)	143
4.11 STUDY 6: THE NEW METAETHICS QUESTIONNAIRE (NMQ).....	149
4.12 STUDY 7: FOLK MORAL OBJECTIVISM SCALE (FMO).....	149
4.13 GENERAL DISCUSSION	150
4.14 CONCLUSION	154
CHAPTER 5: MULTIPLE CHOICE & INDETERMINANCY PARADIGMS.....	156
5.0 INTRODUCTION	156
5.1 MULTIPLE CHOICE PARADIGMS	156
5.2 STUDY 2: INDETERMINACY IN NON-METAETHICAL DOMAINS.....	182
5.3 CONCLUSION.....	214
CHAPTER 6: THIRD PERSON PARADIGM & NORMATIVE ENTANGLEMENT	216
6.0 INTRODUCTION	216
6.1 STUDY 1: THIRD PERSON PARADIGM (PILOT)	229
6.2 STUDY 2: THIRD PERSON PARADIGM WITH CONCRETE MORAL ISSUES.....	243
6.3 STUDY 3: THIRD PERSON PARADIGM (NONMORAL DOMAINS).....	243
6.4 STUDY 4: NORMATIVE ENTANGLEMENT.....	257
6.5 GENERAL DISCUSSION	279
6.6 CONCLUSION.....	289

CHAPTER 7: CONCLUSION	291
REFERENCES	307
APPENDIX A: GLOSSARY	353
APPENDIX B: BIBLIOGRAPHY OF FOLK METAETHICS STUDIES	363
APPENDIX C: METAETHICS SCALES.....	368
APPENDIX D: RESEARCH MATERIALS	375
APPENDIX E: METAETHICS & THE MORAL CONVENTIONAL DISTINCTION	426

PREVIEW

CHAPTER 1:

Introduction

Alex and Sam are two typical Americans. They are planning to have a barbeque tomorrow, but according to a local weather forecast, it might rain. If it does rain, they will have to cancel their plans.

This has led to an argument about the weather:

Alex: “*Given the indeterministic nature of the universe, it is ultimately unknowable whether it will rain tomorrow.*”

Sam: “*Nonsense. It will definitely both rain and not rain tomorrow. The universe will diverge down two separate branches. It will rain in one universe, and not in the other.*”

Alex and Sam seem to be endorsing the Copenhagen and Many Worlds interpretations of quantum mechanics, respectively. This disagreement probably seems absurd. It is incredibly unlikely that interpretations of quantum mechanics would arise in everyday disputes about the weather, and even if they did, neither perspective would be helpful in resolving the practical question of whether to cancel a barbeque. The scenario is also absurd because *most people don't have a position on how to interpret quantum mechanics*. Such considerations are irrelevant to everyday decisions, and nothing about the way ordinary people think or speak requires them to take a side in disputes that occupy theoretical physicists.

I believe that a similar absurdity plagues empirical research on *folk philosophy*. *Folk philosophy* refers to the philosophical stances and commitments of ordinary people. *Philosophical stances* are the philosophical beliefs ordinary people hold. *Philosophical commitments* are the philosophical positions implicit in the way ordinary people speak, think, and act, independent of any particular psychological states.¹ *Ordinary people* are people who lack significant formal philosophical training and have not

¹ I will often refer to these as simply *stances* and *commitments*.

engaged in significant philosophical reflection. The goal of empirical research on folk philosophy is to *describe the philosophical stances and commitments of ordinary people*.

Researchers have devised a cunning array of paradigms for cataloging the stances and commitments of ordinary people, spanning every major branch of philosophy from *metaphysics* (Dink & Rips, 2017; Korman & Carmichael, 2017; Nichols & Bruno, 2010; Paul, 2010; Rose, Schaffer, & Tobia, 2018) to *epistemology* (Nagel, San Juan, & Mar, 2013; Starmans & Friedman, 2012; Swain, Alexander, & Weinberg, 2008; Weinberg, Nichols, & Stich, 2001) to *ethics* (Alicke, & Gordon, & Rose, 2013; Greene, 2008; May & Holton, 2012; Phillips, Nyholm, & Laio, 2014) and *aesthetics* (Cova et al., 2015; Cova & Pain, 2012; Rabb et al., 2020).²

Researchers have frequently interpreted these studies as evidence that ordinary people have stances or commitments that correspond to the categories, accounts, and distinctions recognized by philosophers, e.g. *deontology* and *consequentialism* (e.g. Greene, 2008; Kahane et al., 2018; Johansson-Stenman, 2012; May, 2014; cf. Mihailov, 2022), *compatibilism* and *incompatibilism* (e.g., Carstensen, 2022; Nadelhoffer et al., 2020; Nahmias & Murray, 2011; Nahmias et al., 2004; 2005; 2006; 2015; Nichols, 2012; Nichols & Knobe, 2007, cf. Nadelhoffer, Murray, & Murray, 2021), *psychological* and *non-psychological* conceptions of *personal identity* (e.g. Nichols & Bruno, 2010; Shoemaker & Tobia, 2022; Strohminger & Nichols, 2014; 2015; cf. Starmans & Bloom, 2018a; 2018b), adherence to or rejection of the ‘*ought*’ *implies* ‘*can*’ *principle* (Buckwalter & Turri, 2015;

² Some of these studies may only be concerned with measuring philosophical *intuitions* of a kind that do not correspond to what I mean by philosophical *stances* and *commitments*. “Intuition” is a term regularly employed by philosophers, but for which there is unfortunately no single definition. Researchers studying folk *intuitions* may or may not be studying stances or commitments, e.g., they might be studying *dispositions* to believe certain propositions, rather than studying what people pretheoretically believe are committed to (Earlenbaugh & Molyneux, 2009). Since I am offering a stipulative set of terminological distinctions other authors have not used, I would be unable to confirm whether they conceive of their findings as indicators of stances and commitments without consulting each author, explaining the distinction, and asking them for their position. Failing that, I could selectively opt for studies that appear to fit my distinctions, this task would be largely guesswork on my part. Instead, I chose broadly representative research characterized by scope and impact. I believe these studies provide a more useful picture of the kind of research I am referring to than a curated list of more obscure studies that would risk presenting a skewed notion of what folk philosophical research is about, even if some of these studies are not subject to my criticisms.

Chituc et al., 2016; Cohen, 2018; Henne et al., 2016; 2019; Kissinger-Knox, Aragon, & Mizrahi, 2018; Mizrahi, 2015; Semler & Henne, 2019; cf. Thompson, 2022) *causal-historical* and *descriptivist* theories of reference (e.g. Machery et al., 2004, Mallon et al., 2009; van Dongen et al., 2021), and so on.

Yet few researchers have seriously considered the possibility that ordinary thought and language is *indeterminate* with respect to at least some of these philosophical distinctions.³ In other words, there may be philosophical distinctions that are absent from the way ordinary people speak and think, and that, for these philosophical issues, *ordinary people don't have any particular philosophical stances or commitments at all*. If so, then there may be no way to resolve competing analyses of the philosophical content of folk philosophy, since these analyses could equally accommodate (or fail to accommodate) the data (Gill, 2009). If this is the case, the implications would be catastrophic for much ongoing research, since this research would be attempting to describe features of ordinary thought and language that don't exist.

I am not claiming that ordinary people have *no* philosophical stances or commitments. Even the most skeptical account of folk philosophy would acknowledge that ordinary people hold some minimally construed stances and commitments. To the extent that people endorse first-order moral judgments (e.g., “torture is morally wrong”) they could be said to have a philosophical stance. And insofar as ordinary people speak as though some beliefs are more justified than others, that there is an external world, or that we have a *prima facie* duty to keep our promises, ordinary people may be said to have epistemic, metaphysical, and moral commitments, respectively. I am not challenging the existence of determinate stances and commitments of this kind. Yet these are mundane and uncontroversial beliefs that are rarely the focus of research on folk philosophy.

³ Gill (2009) and Pravato (2020) are notable exceptions. Gill explicitly argues for the possibility of indeterminacy in folk metaethics (for replies see Johansson & Olson, 2015; Sinnott-Armstrong, 2009). Pravato (2020) argues that normative language (e.g., “good,” “ought”) is indeterminate, though it is not meaningless.

Research on folk philosophy instead tends to focus on how ordinary people think about ongoing disputes central to academic philosophy, often with the goal of illuminating or resolving these disputes (Alexander, Mallon, & Weinberg, 2010; Knobe, 2007; Sytsma & Livengood, 2012). As a result, most research on folk philosophy offers insight into ordinary thought and language that is nonobvious or at least subject to reasonable doubt. For instance, researchers studying folk philosophy could plausibly wonder whether ordinary people are committed to the notion that *'ought' implies 'can'*. Yet researchers studying folk philosophy would have little interest in studying whether people think it's morally wrong to torture babies for fun. This is not to say that researchers wouldn't ask this sort of question, but simply that responses to it would not be interpreted as discoveries about folk philosophy, and they wouldn't play an important role in resolving philosophical disputes. Such measures would instead serve some conventional psychological purpose, such as measuring psychopathy.

At the other extreme, it seems unlikely that ordinary people hold philosophical stances about more esoteric philosophical issues, or speak and think in ways that exclusively fit one or another side of obscure philosophical disputes. For example, Gill (2009) finds it implausible that ordinary people speak or think in ways that commit them to mathematical Platonism or anti-Platonism when they use mathematical language, since he believes "[t]he way people use numbers in everyday math simply does not contain answers to the questions that animate philosophy of mathematics" (p. 218). Likewise, it is unlikely that the way people speak or think commits them to a particular stance about the legitimacy of the analytic-synthetic distinction, the optimal decision-theoretic solution to Newcomb's problem, or other debates that concern academic philosophers, and it is even less plausible that ordinary people have explicit stances on these issues.⁴

Thus, there is a continuum between folk stances and commitments that uncontroversially exist and those that are likely absent from folk philosophy altogether. Somewhere between these extremes

⁴ For other examples or popular philosophical disputes, see Bourget and Chalmers (2014; ms).

lies a nebulous middle ground of philosophical accounts, concepts, and distinctions that possess a less-than-certain status as features of folk philosophy. This middle ground is the uncertain territory where I intend to build my case for folk indeterminacy. In making my case, I am not claiming that there are *no* determinate features of folk philosophy between the extremes. There are many features of the way ordinary people think and speak that have been or could be discovered, and at least some of these findings may correspond to traditional philosophical distinctions (e.g., perhaps most ordinary people really are *compatibilists* or *incompatibilists* about “free will”). Yet each hypothesis must be evaluated on an individual basis. We should not presume that a given distinction that is important to philosophers is a part of folk philosophy. In the end, empirical evidence will be the final arbiter of the content of folk philosophy.

While my concerns apply to folk philosophy in general, I will focus exclusively on *folk metaethics*. Metaethics is a branch of philosophy that deals with abstract questions about the nature of morality, such as whether there are moral facts, what makes moral facts true, and how we might acquire knowledge of moral facts (Sayre-McCord, 2014). *Folk metaethics* is simply the subset of *folk philosophy* dedicated to studying the metaethical stances and commitments of ordinary people. While there are several questions that fall under the purview of metaethics, most research focuses on whether ordinary people endorse or are committed to some form of *realism* or *antirealism* about moral facts (Pölzler &

Wright, 2020a; 2020b). Roughly speaking, this distinction concerns whether there are *stance-independent*⁵ facts about what is morally right or wrong (Shafer-Landau, 2003, p. 15).⁶

Realism: *There are stance-independent facts about what is morally right or wrong*

Antirealism: *There are no stance-independent facts about what is morally right or wrong*⁷

Most studies have found that the majority of ordinary people are *metaethical pluralists* who endorse realism about some moral issues and antirealism about others (Beebe, 2014; Beebe et al., 2015; Beebe & Sackris, 2016; Goodwin & Darley, 2008; 2012; Pölzler & Wright, 2019; 2020a; 2020b; Wright, 2015; 2018; Wright, Grandjean, & McWhite, 2013; Zijlstra, 2019; cf. Beebe, 2020). For instance, a person may judge that there is a stance-independent moral fact about whether murder is morally wrong, but that there is no stance-independent moral fact about whether abortion is morally wrong. They might instead think that claims about abortion can only be true or false relative to the moral standards of different individuals or cultures, or that such claims merely expression nonpropositional content, e.g., a negative emotion or an imperative to not get an abortion.

⁵ I follow Shafer-Landau's lead by using the term stance-independent rather than the more common term mind-independent. This is because, as Shafer-Landau notes, the latter is more disposed to prompt confusion since there is one respect in which realism entails that moral facts are mind-independent: they are not made true by the beliefs or values of people. Yet there is another respect in which some moral facts may crucially depend on people's attitudes or values. People sometimes mistakenly think that mind-independence means that whether an action is right or wrong depends on the psychological impact that action would have on an individual. For instance, facts about whether it would be okay to hit someone would depend on whether that person would suffer. And since suffering is a mental state, we might think of claims like "hitting someone is wrong" as mind-dependent, since the reason hitting someone may be wrong is because it causes suffering. Yet this is simply a different sense in which a moral fact could depend on mental states than is intended by the notion of mind-independence or stance-independence.

⁶ Realism is sometimes defined more minimally, and instead consists of two claims: that moral sentences are truth-apt and at least some of them are true (Sayre-McCord, 2015). There seems to have been a shift in recent decades towards reserving the term realism for more robust forms of realism that include stance-independence or other considerations.

⁷ Technical terms like "stance-independent" are likely to prove unhelpful. Roughly, the distinction concerns whether there are moral facts that are not made true by our goals, standards, or values, or whether there are no such facts. Comparison with more familiar concepts might help. Many of us, on reflection, may agree that scientific facts don't depend on our goals, standards, or values. Believing, or really wanting the earth to be flat couldn't make it true that it was flat. Yet many of us would, if we reflected on the matter, deny that there are any stance-independent facts about which food or music is best. We might instead insist that it's a matter of personal preference, and that facts about which food tastes good or bad can only be understood relative to different people's stance towards the food (e.g., whether it tastes good or bad to them), if there are any facts at all. If so, we might say that we're gastronomic antirealists: we deny that there are facts about which foods taste good or bad that are stance-independently true. See Loeb (2003) for an amusing discussion of gastronomic realism and its relation to moral realism.

Some researchers have accepted these findings at face value and argued that *folk metaethical pluralism* is the best interpretation of the data (Davis, 2021; Feltz & Cokely, 2013; Hopster, 2019; Pölzler, 2017; Wright, Grandjean, & McWhite, 2013). Others have interpreted findings in folk metaethics to support the claim that most people are realists (Goodwin & Darley, 2008; 2012) or antirealists (Beebe, 2020; Sarkissian et al., 2011; Pölzler & Wright, 2020b) or at least that we can rule out specific folk metaethical views such as folk noncognitivism (Pölzler & Wright, 2020a). All of these interpretations share the presumption that folk metaethics is *determinate* (Gill, 2009).⁸ In other words, they presume that we can decisively demonstrate that ordinary people adopt a realist or antirealist stance towards moral claims, or speak in ways that best fit some form of realism or antirealism.

I contend that all of this research is fundamentally flawed, and that these interpretations are all mistaken. I don't mean that the studies are poorly designed, and that, with a few tweaks, we'll be able to properly assess what the folk think about realism and antirealism. Nor do I mean that a proper interpretation of these studies would allow us to determine which stances or commitments ordinary people have regarding realism and antirealism. I mean that such efforts cannot succeed in principle because, with few exceptions, *ordinary people have no determinate stances or commitments about the truth status of moral claims*. With respect to realism and antirealism, most ordinary people exist in a state of philosophical superposition that collapses only by engaging in philosophy. I will refer to this as the *metaethical indeterminacy thesis*, which is the claim that *ordinary people have no determinate metaethical stances or commitments with respect to moral realism and antirealism*.

The primary purpose of this dissertation is to explain why I think this is the case, and to convince you that *metaethical indeterminacy* is a plausible account of folk metaethics. Since all existing research on folk metaethics has been interpreted as evidence that ordinary people have determinate

⁸ One exception to this is Pölzler and Wright (2019; 2020a). They don't presume determinacy so much as explicitly argue that their findings cast doubt on indeterminacy. I am indebted to them for clearly laying out their reservations and doubts, and for offering such high quality and insightful studies for our disagreements to center around.

metaethical stances and commitments, this leaves me with three central tasks: (1) to demonstrate that the measures used in existing research on folk metaethics are invalid, (2) to demonstrate this can be readily explained by metaethical indeterminacy and (3) to provide empirical evidence that supports metaethical indeterminacy.

My third objective is intentionally modest. I do not intend to provide anything even approaching decisive evidence of metaethical indeterminacy. This would be an incredibly difficult task, since there are few straightforward ways to convincingly demonstrate that ordinary people have no determinate metaethical stances or commitments. Researchers could always insist that they do, but that existing methods have simply failed to reveal them. Such a task, if it is to succeed, will call for multiple, converging lines of evidence and a broader theoretical foundation than is feasible for the scope of this project. My efforts here should be seen, instead, as laying the groundwork for such a project, providing both a theoretical rationale for the plausibility of metaethical indeterminacy and preliminary empirical support.

It should also quickly become apparent *why* demonstrating metaethical indeterminacy will prove difficult. Although the folk metaethics literature is new and of modest size, no researchers have interpreted their findings as evidence of metaethical indeterminacy. Quite the contrary, all existing research supports a determinate account of folk metaethics, with the only question being what the distribution of different folk metaethical positions is (i.e., whether most people are some form of realist, antirealist, or a combination of the two). Overturning an entire literature is no easy feat. Nevertheless, I *do* believe I can decisively succeed at this task. The methodological critiques that I level against existing studies on folk metaethics reveal far more than the mere difficulty of measuring folk metaethics, they raise serious doubts about whether there is any viable way to measure folk metaethics *at all*. Even if the case for metaethical indeterminacy cannot be decisively established, but only gestured at, I nevertheless hope to show that existing studies fail to provide good evidence that people have

determinate metaethical stances and commitments. In other words, before demonstrating that I am right, I first set out to demonstrate that everyone else is wrong.

If I am correct that ordinary people have no determinate metaethical stances or commitments, the implications would be considerable. Whole research lines in metaethics will prove little more than misguided attempts to measure metaethical stances that most people simply don't have, or to measure commitments that are almost entirely absent from folk philosophy, and, where present, do not vindicate traditional metaethical accounts or existing interpretations of the data. This troubling possibility has been obscured by flawed research design, poor measurement, and a host of methodological problems (Beebe, 2015; Bush & Moss, 2020; Pölzler, 2018b; 2018c). Taken together, these issues give the superficial appearance of nascent fields of research that simply need to refine their tools before progress can be made.

In the chapters that follow, I will address these methodological shortcomings, and argue that researchers have not defended the most plausible interpretation of available data. Yet in doing so, I do not wish to give the impression that the solution is to devise better measures. Rather, I hope to show that the observed patterns of results, as well as the underspecified questions, ubiquitous confounds, pervasive ambiguity in scale items, and other methodological shortcomings are not merely the result of flawed research design, but an unavoidable byproduct of the mismatch between philosophical theories and ordinary thought and language (Bush & Moss, 2020). As I will show, most participants do not understand questions about metaethics in the ways researchers intend, and I will argue that a plausible explanation for this is not that researchers have yet to ask questions in a way participants interpret as intended, but that the types of responses researchers are trying to elicit simply aren't features of folk thought and language to begin with.

I don't believe the problem of folk indeterminacy is a problem exclusive to metaethics. There may be broad ramifications of the kind of analyses I marshal to support this conclusion. If an entire

field of research could present the superficial appearance of legitimate psychological phenomena where none exist, or at least exist in a form radically different from what researchers propose, could the same be true of other lines of research as well? Much to my alarm, I fear an answer in the affirmative.

Given these potentially broader implications, it would be reasonable to ask why I focus on metaethics in particular. Why choose an obscure field with a comparatively small body of literature? The main reason is simply that I'm interested in the topic. But focusing on folk metaethics is also strategic. I am more familiar with research on folk metaethics than any other area of study, and metaethics is my primary area of interest in philosophy. Most importantly, it is an area where the case for indeterminacy is especially strong. Finally, I also suspect that I can build a more persuasive case by focusing on a single topic than by spreading myself too thin by addressing many different areas of research.

If I am successful, this will open the door to future efforts to build similar cases for indeterminacy in other areas of folk philosophy. I am currently developing a similar case for indeterminacy about folk notions of free will and will turn my attention to other folk philosophical distinctions in the future.⁹ But for now, folk metaethics is a rich, active, and evolving field of research that provides fertile ground for developing a template for similar arguments in favor of indeterminacy in other areas of folk philosophy. My critique of research on folk metaethics is thus intended also as a proof of concept for the broader possibility of folk indeterminacy with respect to other philosophical distinctions, and that other areas (or potential areas) of research on folk philosophy may be dead ends.

⁹ I could have presented a similar case for indeterminacy with respect to other areas of folk philosophical research, such as *folk epistemology* (Beebe, 2012; Kim & Yuan, 2015; Machery et al., 2004; Nichols, Stich, & Weinberg, 2003; Weinberg, Nichols, & Stich, 2001) or *personal identity* (De Freitas et al., 2017; Molouki & Bartels, 2017; Shoemaker & Tobia, 2019; Starmans & Bloom, 2018a; 2018b; Strohminger & Nichols, 2014; 2015), among other possibilities. Hopefully researchers who are more familiar with these bodies of literature will investigate the possibility of indeterminacy in these areas of research as well. Even if there are determinate folk views with respect to these philosophical literatures, my approach may aid in refining the measurement tools used to assess other topics.