Doctoral thesis submitted to the Faculty of Behavioural and Cultural Studies Heidelberg University in partial fulfillment of the requirements of the degree of Doctor of Philosophy (Dr. phil.) in Psychology

Title of the publication-based thesis On the Compatibility of Philosophical and Lay Concepts of Moral Uncertainty

presented by Maximilian Johannes Theisen

year of submission 2024

Dean: Prof. Dr. Guido Sprenger Advisor: Prof. Dr. Dr. h.c. Joachim Funke

Abstract

There is an ongoing philosophical debate on the role that uncertainty about normative questions should play in moral decision-making. This debate has been ignited by the surge of *moral uncertaintism*, which describes philosophers' attempts to identify guidelines and heuristics for making better choices in the face of such uncertainty. From a psychological perspective, surprisingly little has been said about the uncertainty that people experience in complex moral cases. In this dissertation, I studied to what extent moral psychology might benefit from recent advancements in the moral uncertaintism literature. In particular, I tested the compatibility between moral uncertaintist and lay views regarding (1) the existence of moral uncertainty due to normative questions, (2) the metaethical interpretation of such normative uncertainty, and (3) normative uncertainty's relevance for moral decision-making. Findings from the first manuscript confirmed the existence of two psychological dimensions that closely track the distinction between empirical and normative uncertainty that is upheld in the philosophical literature. In the second manuscript, discrepancies in the metaethical underpinnings of normative uncertainty between moral uncertaintist and lay concepts became apparent. Finally, the third manuscript revealed that moral ignorance, as the limiting case of normative uncertainty, was not perceived as exculpatory by lay people, which casts doubt on whether lay people think that normative uncertainty should be considered in moral decision-making. I discuss the implications of these findings for the empirical study of uncertainty in moral judgments and decisions. All in all, this work lays the foundation for a philosophically-informed psychology of moral uncertainty.

Acknowledgments

Completing a dissertation provides the pleasant opportunity to look back and thank everyone who has been part of its creation. First, I would like to thank my supervisor, Prof. Joachim Funke, for putting his full trust in me and my work. You gave me freedom to pursue my philosophical interests, for which I am deeply grateful. Likewise, I would like to thank Prof. Werner Greve for agreeing to review this thesis. You always encouraged me to explore the philosophical sides of our discipline which led to this dissertation.

I would like to thank my co-author Markus Germar for his reliable support over the years. Thank you for our many virtual conversations, your strategic guidance, and for ensuring the financial funding of our final study. I am also grateful to James Andow for joining me in the empirical study of moral ignorance. Thank you for your openness and trust to pursue these ideas with a total stranger and for strengthening the philosophical side of this dissertation.

I owe to Prof. Christina Bermeitinger who gave me confidence in my academic skills and kindly supported my application for the dissertational grant. I also owe to Daniel Holt for his constant engagement with my work and the statistical and practical advice he has given me over the years. I would also like to thank the people at the Psychological Institute and the Graduate Academy of Heidelberg University for the continuous support.

I am grateful for the inspiring discussions with participants of the 2021 SoDoc workshop (Leipzig, Germany; online), the 2022 International Conference on Social Dilemmas (Copenhagen, Denmark), the 2022 UK Experimental Philosophy Workshop (Norwich, vi PHILOSOPHICAL AND LAY CONCEPTS OF MORAL UNCERTAINTY

UK; online), the 2023 European Experimental Philosophy Conference (Zurich, Switzerland), and the 2024 research seminar at the Interdisciplinary Centre for Ethics (Kraków, Poland). I am also indebted to the many people who took part in my research. I thank them for sharing their personal moral dilemmas and for giving me highly appreciated feedback on my studies. I would also like to thank Julia Lemp, Daniel Holt, Malvin Escher, Isabell Müller, and James Andow for their helpful comments on various parts of this dissertation.

This project was generously funded by the Baden-Württemberg *Landesgraduiertenförderung* and by project and travel grants from the Graduate Academy of Heidelberg University. I am very thankful for their financial support.

Mein tiefer Dank gilt meinen Eltern – Agnes Blatt-Theisen und Heinz Theisen. Eure Zuversicht und Geduld haben mich durch diese Arbeit getragen. Danke für all das, was ihr mir mit auf den Weg gegeben habt. Mein besonderer Dank gilt Julia Lemp. Danke für das viele Gute in den letzten Jahren.

List of Publications

This dissertation is based on the following publications. The full manuscripts including

supplementary material can be found in the appendix to this work.

Publication 1

Theisen, M. & Germar, M. (in press). Uncertain facts or uncertain values? Testing the distinction between empirical and normative uncertainty in moral judgments. *Cognitive Science*. https://doi.org/10.1111/cogs.13422

Author contributions:

- Theisen: Conceptualization, Methodology, Software, Formal analysis, Investigation, Resources, Data Curation, Writing - Original Draft, Visualization.
- Germar: Conceptualization, Methodology, Resources, Writing Review & Editing, Funding acquisition.

Publication 2

Theisen, M. (2023). Metaethical intuitions in lay concepts of normative uncertainty. *Philosophical Psychology*. Advance online publication. https://doi.org/10.1080/09515089.2023.2174422

Publication 3

Theisen, M. & Andow, J. (2024). *The exculpatory potential of moral ignorance: Evidence from a blame-updating paradigm*. Manuscript submitted for publication.

As of March 15, 2024, this paper is under review at Philosophical Studies.

Author contributions:

- Theisen: Conceptualization, Methodology, Formal analysis, Software, Investigation, Data Curation, Writing – Original Draft, Visualization, Project administration, Funding acquisition.
- Andow: Conceptualization, Methodology, Writing Review & Editing.

Contents

Abstract						
Acknowledgments						
Lis	st of I	Publica	ations	vii		
1	1 Introduction					
2	Theoretical and Empirical Background					
	2.1	Norm	ative Theories of Choice Under Normative Uncertainty	5		
		2.1.1	The Motivation for a Normative Theory of Choice Under Normative Uncertainty	5		
		2.1.2	Formalizations of Choice Under Normative Uncertainty	8		
		2.1.3	An Exemplary Normative Theory of Choice Under Normative Uncertainty	9		
	2.2	Desc	riptive Theories of Choice Under Normative Uncertainty	13		
		2.2.1	The Motivation for a Descriptive Theory of Choice Under Normative Uncertainty	13		
		2.2.2	The Role of Normative Uncertainty in Psychological Paradigms for Studying Moral Conflict	15		
	2.3	The C	Current Research	20		
3	Overview of Manuscripts					
	3.1	.1 Testing the Assumption of an Empirical-Normative Distinction (Manuscript 1)		23		
	3.2	Testir	ng the Realism Assumption (Manuscript 2)	26		
	3.3	Testir	ng the Internalism Assumption (Manuscript 3)	30		
4	Discussion					
	4.1	4.1 Implications for Descriptive Theories of Choice Under Normative Uncertainty				
	4.2	Conc	lusion	40		
References						
Αŗ	pend	dix		53		
	A-1	Manu Betwo	script 1 – Uncertain Facts or Uncertain Values? Testing the Distinction een Empirical and Normative Uncertainty in Moral Judgments	55		
	A-2	Manu	script 2 – Metaethical Intuitions in Lay Concepts of Normative Uncertainty	137		
	A-3	Manu Blam	script 3 – The Exculpatory Potential of Moral Ignorance: Evidence From a e-Updating Paradigm	209		
Declaration in accordance to § 8 (1) c) and d) of the doctoral degree regulation of the Faculty						

1 Introduction

Generally, people can tell right from wrong. They know that it is right to keep promises, to treat others with respect, or to help those in need (Hofmann et al., 2014). Yet, these moral certainties can blur: Our duty to keep a promise can be overridden by another obligation, but it might not always be clear when this is the case (W. D. Ross, 1930, 2002). Similarly, it might not always be clear what treating others with respect specifically entails (Rosenthal, 2019). Last, helping those in need might come at a personal cost, leading to uncertainty about how much we are obliged to bear for others (Hicks, 2018). In short, there can be uncertainty about what morality requires. Such uncertainty is the subject of this dissertation.

I will focus mostly on moral uncertainty tracing back to purely normative questions, such as how to weigh one moral requirement against another. At several occasions, I will also refer to moral uncertainty tracing back to empirical questions, such as what the tangible consequences of a specific action might be. I will use the umbrella term *moral uncertainty* to refer to any doubt about specific moral judgments or decisions and I will use the terms *empirical uncertainty* and *normative uncertainty* to refer to different kinds of moral uncertainty, depending on whether the underlying questions are empirical or normative in nature.¹

¹ In the philosophical literature, various terms have been used to describe this distinction. For example, moral uncertainty due to empirical questions has been called *empirical uncertainty* (MacAskill et al., 2020; Tarsney, 2017), *descriptive uncertainty* (Podgorski, 2020; Robinson, 2022; Robinson & Steele, 2022), *factual uncertainty* (Weatherson, 2019), *non-normative uncertainty* (Sepielli, 2009), or *derived moral uncertainty* (Eriksson & Francén Olinder, 2016). Uncertainty due to purely normative questions has been called *normative uncertainty* (MacAskill, 2014; Podgorski, 2020; Sepielli, 2009; Staffel, 2019), *moral uncertainty* (MacAskill et al., 2020; Tarsney, 2017), or *fundamental moral uncertainty* (Eriksson & Francén Olinder, 2016).

A further distinction becomes relevant when considering the literature on *choice under uncertainty*. *Normative theories* of choice under uncertainty address how people *should* decide from a rational or moral perspective (Hansson, 2013; Steele & Stefánsson, 2020). *Descriptive theories* are concerned with how people *actually* decide (Chandler, 2017). Both types of theories can either refer to choice under empirical uncertainty or to choice under normative uncertainty.

For empirical uncertainty, normative theories describe what would be the most rational (Parfit, 1984; Steele & Stefánsson, 2020) or the most moral way (e.g., Altham, 1983-1984; Doyle & Ericson, 2003; Hansson, 2013; Hayenhjelm & Wolff, 2012; Lewens, 2007; Parfit, 1984; Tarsney, 2018b) to act in the light of uncertainty about the circumstances of a given situation. For example, normative theories might inform us under which conditions it might be appropriate to subject other people to a certain risk (Hansson, 2013). Descriptive accounts of choice under empirical uncertainty address how empirical uncertainty actually affects people's moral judgments and decisions, for example by describing the effect of uncertain outcomes on people's moral evaluation of an action (e.g., Crutchfield et al., 2023; Fleischhut et al., 2017; Kortenkamp & Moore, 2014; Meder et al., 2019; Merlhiot et al., 2018; Ng et al., 2023; Ryazanov et al., 2018; Ryazanov et al., 2021; Shenhav & Greene, 2010; Shou et al., 2020; Shou & Song, 2017; Tversky & Kahnemann, 1981).

For normative uncertainty, normative theories are supposed to describe the rationally or morally right way of taking uncertainty about fundamental normative questions into account when making a moral judgment or decision.² For example, such theories might

² There is dissent as to whether normative models of choice under normative uncertainty make rational or moral prescriptions (see for example Geyer, 2018; Hicks, 2018; Rosenthal, 2021; Sepielli, 2017; Weatherson, 2019). For this dissertation, I will remain agnostic about this issue.

inform us under which conditions normative uncertainty about the moral value of an action is in itself a sufficient reason not to commit the action (MacAskill et al., 2020; Moller, 2011; Rosenthal, 2019). Descriptive models, on the other hand, would focus on how lay people take their normative uncertainty into account when making moral decisions, for example by investigating how they weigh conflicting moral requirements when they are uncertain about their relative importance (Bykvist, 2017; Makins, 2021; Robinson & Steele, 2022).³

In the past few decades, philosophy has seen a surge of interest in normative theories of choice under normative uncertainty (e.g., Bykvist, 2017; Gracely, 1996; Guerrero, 2007; Hicks, 2018; Lockhart, 2000; MacAskill, 2014; MacAskill et al., 2020; Oddie, 1994; Rosenthal, 2019; Sepielli, 2010; Tarsney, 2017; Żuradzki, 2016), including the practical application of such theories to real-world problems (e.g., Bogosian, 2017; Greenhalgh, 2021; Martinho et al., 2021). Descriptive models of normative uncertainty, however, largely remain a desideratum. As an effort to close this gap, I will discuss to what extent the recent philosophical advancements regarding normative theories of choice under normative uncertainty might inform the descriptive study of lay people's choices under normative uncertainty.

The philosophical program that is concerned with the development of normative theories of choice under normative uncertainty has been termed *moral uncertaintism* (Harman, 2015; Sepielli, 2017; for an overview of alternative terms see Robinson, 2022). Models and theories from moral uncertaintism have the advantage that they offer highly formalized accounts of normative uncertainty, which might also be beneficial to the descriptive modeling of such choices (Costa-Gomes & Schoenegger, 2023; MacAskill et

³ I use the term *lay people* to distinguish non-philosophers from expert philosophers.

al., 2020). However, moral uncertaintism relies on non-trivial assumptions about the existence of normative uncertainty, the metaethical interpretation of normative uncertainty, and the relevance of normative uncertainty for moral decision-making. If concepts from moral uncertaintism are to be employed in empirical studies of choice under normative uncertainty, it first needs to be considered whether the underlying assumptions of moral uncertaintism align with moral thought and practice in the addressed lay populations. In this dissertation, I lay the groundwork for the psychological study of normative uncertainty by providing empirical evidence regarding the compatibility of moral uncertaintist and lay concepts of uncertainty in moral contexts.

2 Theoretical and Empirical Background

In section 2.1, I will give an overview of the philosophical project of moral uncertaintism. In section 2.2, I will review how normative uncertainty has been considered in the psychological literature on moral decision-making so far, focusing on cases of moral conflict. Finally, in section 2.3, I will introduce the assumptions underlying moral uncertaintism that are addressed in this dissertation and that may be relevant for the application of moral uncertaintist concepts in descriptive contexts.

2.1 Normative Theories of Choice Under Normative Uncertainty

As a foundation for the discussion of whether moral uncertaintist concepts can be useful for moral psychology, it must first be made clear what the aim of moral uncertaintism is and how moral uncertaintists have attempted to reach it. I will start by illustrating the type of cases that have motivated the moral uncertaintism project (section 2.1.1). Then, I will describe the terms and concepts that are used in moral uncertaintism frameworks to determine the right choice under normative uncertainty (section 2.1.2). Finally, I will give an example of how to apply a normative theory of choice under normative uncertainty to a specific moral case (section 2.1.3).

2.1.1 The Motivation for a Normative Theory of Choice Under Normative Uncertainty

Moral uncertaintism describes the philosophical search for normative theories of choice under normative uncertainty. As such, moral uncertaintists attempt to develop models, rules, and heuristics that help to make better choices in the face of uncertainty about nonempirical questions (Lockhart, 2000; MacAskill et al., 2020; Rosenthal, 2021). To understand the rationale behind this approach, it may prove useful to take a closer look at normative theories of choice under uncertainty in general, before investigating how moral uncertaintists have applied these principles to choice under normative uncertainty.

Normative theories of choice under uncertainty, be it empirical or normative, are concerned with identifying what it would be right for an agent to do given their epistemic state.¹ Epistemic state refers to what an agent knows and believes, including any uncertainty in their beliefs (Harman, 2015, 2022; Kneer, 2018). The sense of rightness that such theories are concerned with is therefore a *subjective* one – it depends on the agent's point of view (Carr, 2015; H. M. Smith, 2010).

That the subjective rightness of an action may depend on the agent's empirical uncertainty is widely accepted among philosophers and non-philosophers (e.g., Harman, 2015, 2022; Kirfel & Hannikainen, 2023; Kissinger-Knox et al., 2018; Monroe & Malle, 2019; Robichaud & Wieland, 2017; Weatherson, 2014, 2019). For example, if a truck driver is uncertain whether there are people in their way or not, it would be subjectively wrong of them to move their truck forward, even if the road is actually clear.²

Moral uncertaintists extend these considerations to normative uncertainty. They argue that the subjective rightness of an action not only depends on the agent's non-moral beliefs, such as whether there are people in front of the truck, but also on their moral beliefs, including their normative uncertainty (e.g., Bykvist, 2017; MacAskill, 2014;

¹ To ensure inclusivity, I will use *singular they* throughout the text, in accordance with APA guidelines (American Psychological Association, 2020).

² This example was inspired by B. Williams' (1981) case of the lorry driver.

Moller, 2011). The claim that moral epistemic states affect subjective rightness has been called *normative internalism* (Tarsney, 2021; Weatherson, 2019).³

As an illustration of how normative uncertainty might affect subjective rightness, consider the case of eating meat (Moller, 2011; Rosenthal, 2019; Weatherson, 2014): A restaurant visitor contemplates whether they should order steak. They are certain that the animal has been brought up in good conditions and that the slaughtering has been conducted according to the current state of the art (we will set aside any social or environmental concerns about meat production for now). Yet, the restaurant visitor is normatively uncertain whether animals have moral status that would make it wrong to treat them as livestock. Would this uncertainty in itself make it subjectively wrong for the restaurant visitor is not fundamentally different from the case of the truck driver. Just as being uncertain about the consequences of moving the truck might in itself give reason not to order steak.

In summary, moral uncertaintists think that empirical and normative uncertainty both affect subjective rightness and should therefore both be considered when making moral decisions. Consequently, as an equivalent to normative theories of choice under empirical uncertainty, moral uncertaintists have attempted to develop normative theories of choice under normative uncertainty that spell out how to identify the best way of moving forward in the face of normative uncertainty.

³ The terms *normative internalism* and *moral uncertaintism* are often used to refer to very similar claims about the relationship between moral epistemic states and subjective rightness (Harman, 2015, 2022; Sepielli, 2017; Tarsney, 2021; Weatherson, 2019). In this dissertation, I use the term *moral uncertaintism* or *moral uncertaintism project* to refer to the search for normative theories of choice under normative uncertainty and I use the term *normative internalism* to refer to the underlying assumption that moral epistemic states, including normative uncertainty, affect subjective rightness.

2.1.2 Formalizations of Choice Under Normative Uncertainty

In their development of normative theories of choice under normative uncertainty, moral uncertaintists have often drawn from concepts and theories about empirical uncertainty (Lockhart, 2000; MacAskill et al., 2020; Żuradzki, 2016). In this section, I will describe how moral uncertaintists have translated the key concepts of choice under empirical uncertainty to the realm of normative uncertainty.

Normative theories about empirical or normative uncertainty usually consider two factors as relevant to what an agent subjectively ought to do: (1) the agent's *certainty* regarding a specific belief and (2) the *conditional value* of the act in question depending on whether the specific belief is true or false (Hansson, 2013; Lockhart, 2000; MacAskill et al., 2020; Steele & Stefánsson, 2020).⁴ The agent's certainty regarding a belief is typically expressed as subjective probabilities ranging from 0% to 100% (MacAskill & Ord, 2020; Savage, 1954). In the case of empirical uncertainty, these probabilities refer to empirical claims, such as the claim that there are people in front of the truck. In the context of normative uncertainty, subjective probabilities refer to normative claims, such as the claim that there are people in front of the truck. In the context of normative uncertainty, subjective probabilities refer to normative claims, such as the

Conditional value describes the objective rightness of an act, depending on whether the respective belief is true or false (MacAskill et al., 2020; MacAskill & Ord, 2020). In the example of empirical uncertainty, the truck driver starting to drive would have a neutral moral value if the road was clear, and it would have a negative moral value if there were

⁴ Some have also argued that the right way to act under normative uncertainty is simply to follow the normative view one is most certain of, which neglects the aspect of conditional value (Gracely, 1996; Gustafsson & Torpman, 2014). In the most recent discourse, however, these views have largely been abandoned, including by some of their former proponents (Gustafsson, 2022).

⁵ Sepielli (2017) speaks of epistemic rather than subjective probabilities in the context of moral uncertaintism, referring to the difference between an agent's moral beliefs and the moral beliefs that it would be epistemically justified for the agent to have. This distinction is relevant for discussions of normative internalism but it will not be of further relevance to this dissertation.

people in front of the truck. In the example of normative uncertainty, ordering the steak would have a neutral moral value if animals do not have moral status and a negative moral value if animals do have moral status.

In conclusion, normative models of choice under normative uncertainty operate with two central concepts: the (subjective) probability of a normative claim being true or false and the moral value of the act in question depending on the respective normative claim. In the following section, I will illustrate a common way in which moral uncertaintists have used these concepts to identify the right choice under normative uncertainty.

2.1.3 An Exemplary Normative Theory of Choice Under Normative Uncertainty

The aim of moral uncertaintism is to provide action guidance for normatively uncertain agents, allowing them to make better choices in light of their moral doubts (Sepielli, 2009; Tarsney, 2021). To show how moral uncertaintist models help to achieve this aim, I will demonstrate how to apply such a model to a concrete moral case.

The theory that has been considered the default view in moral uncertaintism is *maximizing expected moral value* (MacAskill & Ord, 2020). Maximizing expected moral value applies the decision theoretic idea of *maximizing expected utility* to the context of uncertainty about normative claims (MacAskill et al., 2020; MacAskill & Ord, 2020; Steele & Stefánsson, 2020). I will describe how expected moral value is calculated (section 2.1.3.1) and how maximizing expected moral value can be applied to a case about abortion (2.1.3.2).

2.1.3.1 Calculating Expected Moral Value

To calculate the expected moral value of an option, the product of an agent's subjective probability regarding a normative claim and the conditional moral value of the respective

option according to that normative claim is summed up over all normative claims that the agent has at least some credence in (MacAskill et al., 2020). Calculating expected moral value thus not only requires that the agent's subjective probabilities can be estimated (see also MacAskill & Ord, 2020), but also that conditional moral value can be represented numerically and on a common scale across the considered normative claims. These requirements are far from trivial. In particular, comparing conditional moral value across different normative frameworks – so-called intertheoretic value comparison – has been discussed as one of the main challenges of moral uncertaintism (Carr, 2020; Côté, 2023; Nissan-Rozen, 2015; J. Ross, 2006; Tarsney, 2018a; Weatherson, 2019). To illustrate the idea of maximizing expected moral value, I will therefore refer to a simplified case of normative uncertainty where both normative claims in question imply a common scale of moral value.

A further simplification is the assumption of discrete normative claims. In reality, people might not be normatively uncertain about discrete normative claims, such as whether animals have moral status or not, but about a continuous value, such as *how much* moral status non-human animals have in comparison to humans, for example (Bykvist, 2017; DeGrazia, 2008; Rosenthal, 2019). In such cases, subjective probability would not be described as an individual percentage but as a probability distribution across a range of possible values. For illustrative purposes, the following example will be built on the assumption of two discrete normative claims. However, maximizing expected moral value can, in principle, also be applied to uncertainty about continuous values. We will leave the assumption of discrete normative claims behind when discussing normative uncertainty in the context of moral conflict in section 2.2.

2.1.3.2 Maximizing Expected Moral Value in the Context of Abortion Consider the case of a gynecologist who is uncertain whether they should perform an abortion (Moller, 2011; Weatherson, 2014; Żuradzki, 2016). Their patient is 12 weeks pregnant and would prefer not to have a child right now. The gynecologist mostly thinks that their patient's fetus does not have moral status that would make it impermissible to perform this abortion. However, they have some normative uncertainty about this.

To find out what would be subjectively right for the gynecologist to do according to maximizing expected moral value, the expected moral value of the available options must be compared (MacAskill et al., 2020; MacAskill & Ord, 2020). To determine the expected moral value of each option, it first needs to be determined how the gynecologist's subjective probabilities are distributed across the competing normative claims that the fetus does or does not have moral status. Since the gynecologist tends toward the view that the fetus does not have moral status, we might describe them as being 70% certain that the fetus does not have moral status and as 30% certain that it does.

In a second step, it needs to be considered what the conditional moral value of performing or not performing the abortion would be, depending on whether the fetus does or does not have moral status. We assume that it would be very wrong to perform the abortion if the fetus has moral status. Arbitrarily, we might describe this conditional moral value as $-100.^{6}$ We might further assume that it would also be wrong not to perform the abortion if the fetus does not have moral status, since this would deny the patient autonomy over

⁶ Such an arbitrary assignment of conditional moral value is sufficient for the illustration of maximizing expected moral value. When applying maximizing expected moral value to an actual case, however, these values need to accurately represent what is morally at stake according to each normative view that is considered. As mentioned earlier, this might be one of the main problems of moral uncertaintism. For recent attempts to circumvent the problem of intertheoretic value comparison, see for example Carr (2020), MacAskill (2016), Newberry and Ord (2021), or Tarsney (2019).

their reproductive choices. However, this might not be as grave as performing the abortion of a fetus with moral status. Therefore, we might describe this conditional moral value as -40. Last, there are two morally neutral options: Performing the abortion if the fetus does not have moral status and not performing the abortion if the fetus does have moral status. We assign these options a conditional moral value of 0.

The comparison of conditional moral value across both available options shows that, in this exemplary case, not performing the abortion has a better (i.e., less negative) expected moral value than performing the abortion (Table 2.1). Therefore, in this particular case, maximizing expected moral value would require that the gynecologist does not perform the abortion, even though they give most credence to the normative view according to which performing the abortion would be permissible (MacAskill et al., 2020; MacAskill & Ord, 2020).

Option	Fetus does not have moral status (70%)	Fetus does have moral status (30%)	Expected moral value
Performing the abortion	Neutral (0)	Wrong (-100)	70% * 0 + 30% * (-100) = -30
Not performing the abortion	Wrong (-40)	Neutral (0)	70% * (-40) + 30% * 0 = -28

Table 2.1. Calculation of Expected Moral Value for a Case of Abortion

It should be noted that maximizing expected moral value does not necessarily lead to the conclusion that abortions should not be performed. Slight changes in subjective probabilities or in conditional moral value might be sufficient to tip the scale in favor of performing the abortion. For example, it would have been sufficient to describe the conditional moral value of not performing the abortion if the fetus does not have moral status as -50 to yield a better expected moral value for performing rather than not performing the abortion (see also Moller, 2011).

The case of the gynecologist illustrates how moral uncertaintism can provide actionguidance for normatively uncertain agents, depending on how certain the agent is and what is morally at stake. It also became apparent the non-trivial assumptions that must be met to apply such models to actual moral cases, such as the need to describe moral value across various normative views on a common, cardinal scale. In the following chapter, I will turn my attention to the psychological study of choice under normative uncertainty to work out to what extent concepts and models from moral uncertaintism, such as those whose application I illustrated for the gynecologist case, might be suited to inform the empirical study of choice under normative uncertainty.

2.2 Descriptive Theories of Choice Under Normative Uncertainty

Choice under normative uncertainty may not only be a problem for philosophers but also for lay people. How do lay people consider their uncertainty about relevant normative claims when making moral decisions? And do they think that there are better and worse ways to do so? These are psychological questions that need to be addressed empirically.

I will argue that there is a need for a descriptive theory of choice under normative uncertainty given the prevalence of normative uncertainty in everyday life, particularly in the form of moral conflict (section 2.2.1). Then, I will raise concerns about the suitability of current psychological paradigms to address questions about normative uncertainty (section 2.2.2).

2.2.1 The Motivation for a Descriptive Theory of Choice Under Normative Uncertainty

There are a vast number of popular media formats, such as newspaper columns or call-in radio shows, where lay people seek advice about personal moral questions (e.g., Domian,

2008; New York Times Magazine, 1999; Süddeutsche Zeitung Magazin, 2006). The popularity and longevity of such formats – some of them span multiple decades – suggest that difficult moral questions are part of everyday life. Contributors to such formats report being uncertain about moral duties toward abusive parents, about fairness between partners with unequal income, or about moral justifications for breaking the law. Given the complexity of everyday morality that is reflected in these reports, it can be expected that not only philosophers but also lay people regularly experience uncertainty about normative claims.

As mentioned earlier, normative uncertainty in everyday life might not be so much about discrete normative claims, such as whether a specific being has moral status or whether utilitarianism is the right ethical theory (MacAskill et al., 2020; Moller, 2011), but more about claims about a continuous value, such as how important one moral requirement is relative to another (Makins, 2021; McConnell, 2022; Parfit, 2011; Robinson & Steele, 2022; W. D. Ross, 1930, 2002). This latter form of normative uncertainty can appear in moral conflicts.

Moral conflict describes cases in which an agent has moral reason to do each of multiple actions but cannot do all of them (Mason, 1996; McConnell, 2022; Sinnott-Armstrong, 1988). In some cases, moral conflict might be low, so that one moral requirement clearly overrides the others and the conflict is easily resolved. In other cases, however, when moral conflict is high, there might be normative uncertainty about which moral requirement is more important and, consequently, about what would be the right choice to make (*epistemic moral dilemma*; McConnell, 2022).

Empirical studies on the metacognition involved in moral conflicts have confirmed that judging such cases is perceived as difficult (Bai et al., 2021; Behnke et al., 2020; Chen et

al., 2022; Kuhnert et al., 2017) and that people have low confidence regarding their judgments of such cases (Lee et al., 2018; Ryazanov et al., 2021; Shivnekar & Srivastava, 2023). While some of this uncertainty might be due to empirical questions (Shou & Song, 2017), positive correlations between perceived moral conflict and perceived decisional difficulty suggest that the uncertainty associated with such cases is at least partly due to normative rather than empirical questions (Carmona-Perera et al., 2015; Conway & Gawronski, 2013). This interpretation is further supported by findings that resolving conflict between moral values is perceived as more difficult than resolving conflict between non-moral values (Hanselmann & Tanner, 2008).

Both anecdotal and empirical evidence suggest that lay people are familiar with the experience of normative uncertainty, which has mostly been shown for cases of moral conflict. In the following section, I will review the paradigms that have been recently used to study choices in moral conflict situations and whether these paradigms are suited to investigate the role that normative uncertainty plays in such cases.

2.2.2 The Role of Normative Uncertainty in Psychological Paradigms for Studying Moral Conflict

Moral conflict has been one of the central topics of contemporary moral psychology, mostly in the context of sacrificial dilemmas (Ellemers et al., 2019; Greene et al., 2001).⁷ Sacrificial dilemmas are cases where a sacrifice is required to prevent an even larger loss of the same type, for example, sacrificing the life of one person to save the lives of many others (Foot, 1967). In sacrificial dilemmas, obligation-based (deontological) moral

⁷ While sacrificial dilemmas have prevailed in the recent literature, it should be noted that not all moral conflict research has relied on such cases (see for example Hanselmann & Tanner, 2008; Mandel & Vartanian, 2008; Shaddy et al., 2021; Shamoun & Svenson, 2002). For the most part, my arguments regarding the role of normative uncertainty in studies of sacrificial dilemmas can also be applied to other studies of moral conflict.

reasons conflict with consequence-based (utilitarian) moral reasons (Conway & Gawronski, 2013).

When sacrificial dilemmas were first introduced to the psychological literature, they were mostly discussed in light of the *dual-process theory* of moral judgment (Ellemers et al., 2019; Greene & Haidt, 2002; Greene et al., 2004; Greene et al., 2001). Moral judgment in sacrificial dilemmas is here understood as the result of two competing, cognitive subsystems. Processes associated with System 1 are fast, intuitive, and unconscious, while those associated with System 2 are slow, deliberate, and conscious (Evans, 2008; Evans & Stanovich, 2013). Initially, it was argued that deontological inclinations are the result of System 1 and that utilitarian inclinations are the result of System 2 (Greene et al., 2001). More recent findings, however, suggest more flexible associations, with either subsystem being capable of producing responses that either correspond to deontology or utilitarianism (Bago & De Neys, 2019; Białek & De Neys, 2016, 2017). These later findings have led to an increased focus on the conflicting moral principles that lie at the core of sacrificial dilemmas (Skovgaard-Olsen & Klauer, 2024).

The incompatible requirements of deontology and utilitarianism can lead to high moral conflict (Białek & De Neys, 2016, 2017; Mata, 2019; Rosas et al., 2019; Skovgaard-Olsen & Klauer, 2024). Since the judgment of such dilemmas is often perceived to be difficult, it can be expected that people are normatively uncertain how much weight should be given to each of the competing arguments in such situations (Bai et al., 2021; Behnke et al., 2020; Bykvist, 2017; Carmona-Perera et al., 2015; Chen et al., 2022; Conway & Gawronski, 2013; Kuhnert et al., 2017). What have psychological paradigms to say about the role of such normative uncertainty in the judgment of sacrificial dilemmas?

There have been two recent advancements in the study of sacrificial dilemmas that need to be discussed in this context. The first is the application of process-dissociation models (Conway & Gawronski, 2013; Gawronski et al., 2017). The second is the postulate of a moral trade-off system (Guzmán et al., 2022).

2.2.2.1 Process-Dissociation Models

In earlier studies on sacrificial dilemmas, it was not possible to tell whether a participant chose an option out of endorsement of the respective normative view or out of rejection of the opposite normative view, since responses were usually measured on a single scale (Conway & Gawronski, 2013). To address this limitation, Conway and Gawronski (2013) propose to measure participants' decisions across a wide range of sacrificial dilemma-type cases that systematically vary in whether they create moral conflict or not. By using process-dissociation approaches (Jacoby, 1991), it is then possible to separately estimate participants' moral sensitivity to deontological and utilitarian concerns, respectively. Over the past few years, this approach has been further adapted to consider sensitivity to inaction (Gawronski et al., 2017), to estimate parameters on a participant-level rather than a group-level (Körner et al., 2020), and, most recently, to estimate participant's sensitivity to moral conflict itself (Skovgaard-Olsen & Klauer, 2024).

Moral sensitivity to deontological or utilitarian reasons somewhat resembles the idea of subjective probabilities regarding normative claims, as I have discussed it in the context of moral uncertaintism. Yet, there is a central difference between both constructs. While subjective probabilities describe an agent's certainty, moral sensitivity describes how much weight a participant gives to obligations or consequences in the presented scenarios. This weighting might already be the result of a process that considers the agent's beliefs regarding the moral importance of the respective aspect and their

uncertainty regarding those beliefs. Due to the inability to differentiate between certainty and importance, current process-dissociation approaches do not inform us about how people weigh normative uncertainty when judging sacrificial dilemmas.

2.2.2.2 Moral Trade-Off Model

Another perspective on choice in sacrificial dilemmas has recently been proposed by Guzmán et al. (2022). They argue that dual-process theories of moral judgment do not sufficiently account for people's capacity to make compromises between conflicting normative views (but see Greene, 2023). As an alternative, Guzmán et al. describe a cognitive moral trade-off system that allows people to flexibly weigh competing moral requirements in cases of moral conflict.

To show lay people's capacity to make moral compromises, Guzmán et al. used various sacrificial dilemmas. In addition to offering a choice between deontological or utilitarian actions, they also offered several intermediate actions to choose from that were neither purely deontological nor purely utilitarian. Additionally, the researchers systematically manipulated the conditional moral value of the presented options, for example, by varying how many people would be saved or sacrificed. Using this paradigm, Guzmán et al. could show that people often prefer moral compromise over extreme judgments, that people's judgments reflect rational principles, and that people flexibly adapt their choices to the specific moral stakes of the presented scenarios (Guzmán et al., 2022).

The model and the findings presented by Guzmán et al. are relevant for our discussion of normative uncertainty in multiple ways. First, similar to the moral sensitivities described in process-dissociation approaches, Guzmán et al.'s model assumes that people assign weights to competing moral reasons. In contrast to moral sensitivities, however, which are modeled on the person-level, the weights in the moral trade-off model are understood to be specific to the respective moral case. Neither of these concepts, however, differentiate between certainty regarding the respective moral reasons and the importance of these reasons.

Second, Guzmán et al. confirmed earlier findings that people are sensitive toward manipulations of conditional moral value by varying how strong the deontological and utilitarian reasons in the respective cases are (Awad et al., 2018; Bai et al., 2021). This directly corresponds to the moral uncertaintist idea of conditional moral value, and similar manipulations might be used to assess how normative uncertainty affects moral choices.

Last, people's preference for moral compromise might in itself be relevant to discussions of normative uncertainty. There are two possible explanations for choosing moral compromise: Either people are convinced that moral compromise is the best available moral option, or they choose compromise because they are uncertain which of the extreme judgments is the best available moral option. While the first would resemble some form of threshold deontology (Brennan, 1995; Moore, 2010; Rosenthal, 2019), the latter would more resemble moral uncertaintist ideas such as *maximizing expected moral value*, with moral compromise being chosen not because it is the objectively right moral view but because it maximizes expected moral value, for example (MacAskill, 2014; MacAskill et al., 2020; Newberry & Ord, 2021; see also Shaddy et al., 2021; for a discussion of the relationship between threshold deontology and moral uncertaintism see Rosenthal, 2019). Since normative uncertainty is not explicitly considered in Guzmán et al.'s work, their paradigm is not suited to differentiate between these two motivations behind moral compromise.

In summary, while the ubiquity of difficult moral cases suggests that psychology should be concerned with lay people's consideration of normative uncertainty in moral decisionmaking, the paradigms that currently prevail in moral psychology are not suited to address such questions, mostly because normative uncertainty is not considered separately from the importance that people assign to the respective moral reasons. Concepts from moral uncertaintism, such as the separate consideration of subjective probabilities and conditional moral value, might further illuminate how moral sensitivities in process-dissociation models and weights in the moral trade-off model are to be interpreted psychologically.

2.3 The Current Research

As the previous chapter has shown, despite moral psychology's focus on moral conflict, there still is a lack of descriptive theories of choice under normative uncertainty. Concepts from the highly formalized normative models of choice under normative uncertainty, as they have been proposed by moral uncertaintists, might potentially be suited to inform such descriptive theories. Additionally, moral uncertaintism inspires novel research questions, such as which strategy lay people think is the most rational or most moral for choice under normative uncertainty (for first empirical investigations of somewhat related questions, see Spälti et al., 2019).

Before empirical researchers employ ideas from moral uncertaintism in their work, it needs to be ensured that the core assumptions underlying moral uncertaintism are not in conflict with moral thought and practice in the targeted populations. This is important, since moral uncertaintism rests on specific assumptions about moral thought and practice that are rejected by some philosophers and might also not be shared by lay people (Bourget & Chalmers, 2023; Harman, 2015, 2022; Hedden, 2016; Nissan-Rozen, 2015;

Weatherson, 2014, 2019). If there is a discrepancy between lay and uncertaintist concepts of choice under normative uncertainty, this would undermine the benefits of implementing such concepts in empirical work.

For this dissertation, I investigated lay people's views on three core assumptions of moral uncertaintism. The first assumption concerns the existence of normative uncertainty itself. Whether this assumption is met determines the need for a descriptive theory of choice under normative uncertainty. The second assumption concerns the metaethical interpretation of normative uncertainty. This assumption affects which terms and concepts relating to normative uncertainty should be used in its empirical study. The third assumption refers to the normative internalist idea that an agent's moral beliefs affect subjective rightness. This assumption is relevant for the study of lay people's intuitions about the best strategy for choice under normative uncertainty. The assumptions are the following:

- (1) The assumption of an empirical-normative distinction in moral uncertainty: Moral uncertaintists assume that people's moral uncertainty does not always trace back to empirical questions alone but can also trace back to normative questions (Bykvist, 2017; MacAskill et al., 2020). In the first manuscript, my co-author Markus Germar and I tested whether normative uncertainty in lay people exists and whether it is psychologically distinct from uncertainty about morally-relevant empirical facts.
- (2) The realism assumption: The moral uncertaintism project has been linked to the metaethical notion of realism, i.e., the notion that there are moral truths that are independent of moral thought and practice (Brink, 1989; Railton, 1986; Sepielli, 2017; Shafer-Landau, 2003). This is reflected in the terminology and concepts of

moral uncertaintism models, such as describing normative uncertainty in terms of probabilities of normative claims or the assumed symmetry between empirical and normative uncertainty. Yet, lay people have been shown to hold mostly anti-realist intuitions (Goodwin & Darley, 2008, 2012; Pölzler & Wright, 2020). In the second manuscript, I tested whether this discrepancy would lead to a mismatch between lay and uncertaintist concepts of normative uncertainty.

(3) The internalism assumption: Moral uncertaintism rests on the assumption that an agent's moral epistemic state, such as their moral beliefs or their normative uncertainty, matters to what they subjectively ought to do (Harman, 2015, 2022; Robinson, 2022; Sepielli, 2017; Tarsney, 2021). Critics of moral uncertaintism have argued that normative internalism implies exculpation via moral ignorance and have therefore rejected the moral uncertaintism project as misguided (Harman, 2015, 2022). In the third manuscript, my co-author James Andow and I tested under which conditions people would see moral ignorance as exculpatory, thereby testing the intuitiveness of one of the central arguments that have been levelled against moral uncertaintism.

3 Overview of Manuscripts

In the following, I will give an overview of the research that constitutes this dissertation. Findings from Manuscript 1 regarding the distinction between empirical and normative uncertainty in lay people are described in section 3.1. In section 3.2, I report findings from Manuscript 2 regarding the role metaethical intuitions play in lay concepts of normative uncertainty. Finally, in section 3.3, I focus on findings from Manuscript 3 regarding lay views on the exculpatory potential of moral ignorance.

3.1 Testing the Assumption of an Empirical-Normative Distinction (Manuscript 1)

The philosophical literature suggests that people's moral uncertainty sometimes traces back to normative rather than empirical questions (e.g., Harman, 2015; Lockhart, 2000; MacAskill et al., 2020; Nissan-Rozen, 2015). Yet, the existence of normative uncertainty as a psychological state that is distinct from empirical uncertainty has so far not been tested. While psychologists have given some attention to general moral uncertainty in the past few years (e.g., Alsaad et al., 2022; Brannon et al., 2019; Mata, 2019; Poliquin, 2010; Vega et al., 2020), normative uncertainty has yet to be established as a psychological construct. The first manuscript describes the development of the *Moral Uncertainty Scale* (MUS) which introduces empirical and normative uncertainty as two conceptually and empirically distinct dimensions of uncertainty in moral judgments.

The idea behind the MUS was to develop a scale that would measure the degree to which people's uncertainty regarding a specific moral issue traces back to empirical or normative questions. The initial item pool was based on literature on the ethics of risk and moral uncertaintism. The validity of these items was rated by N = 9 expert philosophers

who had published peer-reviewed work on normative uncertainty before. Subsequent cognitive interviews (think-aloud protocols) with N = 10 lay people ensured that the items would be interpreted correctly by members of the addressed population.

Before the developed items were subjected to factor analysis, it had to be decided on the moral scenario the data collection would be based on. The use of a single scenario might lead to the identification of a factor structure that is unique to that scenario and that does not generalize across other moral issues. However, the aim was that the MUS would be applicable to a wide variety of moral issues that are of interest to researchers. To address this challenge, we developed a novel procedure (pre-tested on N = 50 participants) that would prompt each participant to report an individual case of moral uncertainty (see also Gerpott et al., 2018). By using this procedure for data collection, the identified factor structure would not be tied to a specific moral scenario but would be derived from a wide range of moral issues, with each participant responding to the MUS items in the context of a different scenario.

We used this procedure to collect data for exploratory factor analysis. After the removal of items with low factor loadings and of factors with low reliability, data from N = 265 participants suggested three reliable factors of moral uncertainty: "lack of information" and "unclear consequences" described uncertainty about empirical questions, and "normative uncertainty" described uncertainty about normative questions. This lent first suggestive support to our assumption of empirical and normative uncertainty as distinct psychological states.

The identified factor structure was then confirmed and validated on a separate dataset using both factor-analytical and experimental approaches. Here, we switched from participant-generated scenarios to researcher-generated scenarios, as they are typically used in moral psychology (Ellemers et al., 2019). If the factor structure that was derived from participant-generated scenarios could be confirmed using researcher-generated scenarios, this would be highly indicative of the robustness of the MUS factors across a wide variety of moral issues.

For this pre-registered study, we collected data from N = 402 participants who each responded to eight different moral scenarios. The content of these scenarios was subject to experimental manipulations of the certainty of the presented empirical information (low vs. high), the degree of inherent value conflict (low vs. high), and the overall setting (oil ship vs. company). For each scenario, participants responded to the MUS items, items capturing overall moral certainty (O'Connor, 1995), and items capturing attitude clarity and attitude correctness, which are two established dimensions of attitude certainty (Petrocelli et al., 2007).

As expected, the factor structure of the MUS with two empirical and one normative subfactor was confirmed. The MUS exhibited strict measurement invariance across different levels of the experimental manipulations, its subfactors were highly negatively correlated with overall moral certainty (convergent validity), and its subfactors were empirically distinct from attitude clarity and attitude correctness (divergent validity). The experimental manipulations further revealed that scores on the empirical uncertainty subfactors were sensitive to the introduction of uncertain empirical information and that scores on the normative uncertainty subfactor were sensitive to the introduction of value conflict. However, to a lesser degree, each manipulation also affected the subscale that it was not intended to target. These effects could be explained by each subscale not only capturing the *kind* of moral uncertainty but also the *degree* or strength of the uncertainty.

After controlling for overall moral certainty, the specificity of each manipulation was increased.

Findings from the first manuscript showed that the distinction between empirical and normative kinds of moral uncertainty reflects a psychological reality. Using the MUS, empirical and normative uncertainty can be measured separately and reliably across a wide range of moral issues. However, the findings also suggested that the MUS subscales not only capture the type of uncertainty but also the overall degree of uncertainty, which should be controlled for in future research. Overall, the validity of normative uncertainty as a psychological construct and the importance of considering normative uncertainty in the empirical study of moral judgments could be stated.

3.2 Testing the Realism Assumption (Manuscript 2)

The concept of empirical uncertainty closely aligns with the general usage of the term *uncertainty* in non-moral contexts (Hansson, 2013; Steele & Stefánsson, 2020). It describes the state of not knowing with certainty whether an empirical claim is true or false. For normative uncertainty, matters are more intricate. What exactly does it mean to be uncertain about a normative question? Can a normative claim be true or false, as moral uncertaintists usually assume? And if so, what determines its truth-value? These considerations were the starting point of the second manuscript.

Questions about the nature of moral evaluations fall into the realm of metaethics. For the second manuscript, I considered various metaethical notions and their implications for normative uncertainty. By analyzing the metaethical assumptions underlying the moral uncertaintism discourse and the literature on lay people's metaethical intuitions, I identified a potential discrepancy between lay and moral uncertaintist understandings of
morality. Yet, it was unclear whether the discrepancy between lay people's and moral uncertaintists' metaethical assumptions would carry through to a discrepancy in concepts of normative uncertainty. This was tested in two studies by assessing the relationship between metaethical intuitions and concepts of normative uncertainty in lay people.

		Cognitivism		Non-cognitivism
Criterion	Realism	Cultural relativism	Individual relativism	
Understanding of moral judgments	Beliefs about universal, objective facts	Beliefs about facts that can vary between cultures	Beliefs about facts that can vary within cultures	Non-cognitive states (desires, intentions,)
Interpretation of normative uncertainty	Uncertainty regarding universal, objective moral facts	Uncertainty regarding culturally accepted moral values	Uncertainty regarding own moral values	Ambivalent non- cognitive states

Table 3.1. Overview of Metaethical Notions and the Assumed Psychological States of Normative Uncertainty

Note. Reproduced from Theisen (2023).

Table 3.1 gives an overview of the metaethical notions that were considered in these studies. The distinction between cognitivism and non-cognitivism refers to the question of whether moral claims can be true or false. Cognitivism states that moral judgments are expressions of beliefs and can, as such, be true or false. In this light, normative uncertainty is interpreted as uncertainty about whether a specific normative claim accurately reflects the respective moral facts. Non-cognitivism, on the other hand, understands moral judgments as expressions of non-cognitive attitudes, such as desires or intentions. These are typically understood as not truth-apt – they can neither be true nor false (Makins, 2021). How normative uncertainty is to be understood under non-cognitivism is still an open debate, and some have claimed that both concepts are mutually incompatible (MacAskill et al., 2020; M. Smith, 2002). Recently, however, it has been proposed to understand normative uncertainty from a non-cognitivist perspective as conflicting non-cognitive attitudes, similar to ambivalence or polytely (Makins, 2021;

Robinson & Steele, 2022). This was the interpretation of non-cognitivist uncertainty that I followed for the empirical studies of the second manuscript.

Distinctions within cognitivism reflect different interpretations of what constitutes the truth of a moral claim. Here, the focus is on whether moral facts are objective (realism) or whether they depend on human thought and practice (relativism). From a realist perspective, normative uncertainty is uncertainty about what the objective moral facts in a given situation are, independent of how people think about those facts (Brink, 1989). This is the view that has prevailed in the moral uncertaintism discourse (Tarsney, 2021). From a cultural relativist perspective, normative uncertainty refers to culturally defined norms (Tilley, 2000). From an individual relativist perspective, normative uncertainty refers to personal norms and values (Moller, 2011).

In two studies, I measured participants' metaethical intuitions and different aspects of their normative uncertainty concepts for a series of concrete issues. The measurement of metaethical intuitions largely followed tasks presented in Wagner et al. (2021). To capture lay concepts of normative uncertainty, I developed novel tasks that focused on participants' interpretations of normative uncertainty and their views on strategy selection for resolving the uncertainty.

In both studies, participants' concepts of normative uncertainty were linked to their metaethical intuitions and participants indicated mostly anti-realist interpretations of normative uncertainty. In Study 1, most participants interpreted normative uncertainty as ambivalence, followed by relativist notions of uncertainty. There was a moderate association between these interpretations and the participants' metaethical intuitions, which was driven by elevated rates of realist interpretations from realist participants. Realist participants also showed elevated rates of indicating that they would try to resolve

their uncertainty by consulting science or philosophy, while the majority of participants instead chose to increase clarity regarding their own view.

In pre-registered Study 2, the previous findings were replicated. The uncertainty task was updated to sharpen the contrast between cognitivist and non-cognitivist responses, which drove down the number of people interpreting uncertainty as ambivalence. In this updated task, uncertainty was interpreted in relativist ways in more than two thirds of the cases. The associations with the metaethical tasks remained intact. The resolving task now focused on what the aim of resolving the uncertainty would be. Consistent with findings from Study 1, the most frequent response was that the aim was to learn about oneself and one's own views. In the strategy task, which captured how participants would try to reach the respective aim of resolving uncertainty, most participants chose introspection or expert advice. Responses to each of the tasks referring to normative uncertainty were significantly associated with responses to the metaethical tasks.

In summary, the main argument regarding the mismatch between uncertaintist and lay concepts of normative uncertainty was supported by findings from both studies. Being normatively uncertain was mostly not seen as uncertainty of an objective moral fact and some even rejected the idea that normative uncertainty refers to any fact at all. Some of the variance in interpretations of normative uncertainty was explained by participants' metaethical intuitions, which highlights the importance of considering such intuitions when conducting empirical research on normative uncertainty. In most parts, the moral uncertaintist assumption of moral realism was not shared by lay people.

3.3 Testing the Internalism Assumption (Manuscript 3)

In the third manuscript, the focus was on the assumption of normative internalism, i.e., the claim that an agent's moral epistemic state matters to what they subjectively ought to do, just as their non-moral epistemic state does (Geyer, 2018; Harman, 2015; Robinson, 2022; Sepielli, 2017; Tarsney, 2021; Weatherson, 2014, 2019).

One central argument that has been levelled against internalism focuses on cases of moral ignorance, i.e., cases where the agent is certain of a false moral view (Harman, 2015, 2022; Weatherson, 2014, 2019). Moral ignorance is understood as the limiting case of normative uncertainty, where full credence is given to a false normative claim and zero credence is given to the right normative claim. According to internalism, a morally ignorant agent subjectively ought to act in line with their false moral assumptions (MacAskill & Ord, 2020; Weatherson, 2019). Critics argue that, since an agent cannot be blamed for doing what is subjectively right, internalism thus implies exculpation via moral ignorance (Geyer, 2018). Yet, so the critics, since moral ignorance does in fact not exculpate, normative internalism must be false (Harman, 2015, 2022).

Different responses have been raised against this argument. One route that has been taken is the rejection of the premise that an agent who does what is subjectively right cannot be morally culpable (Geyer, 2018; Sepielli, 2017). Another route has been the rejection of the premise that moral ignorance is not exculpatory. Both Geyer (2018) and Sepielli (2017) argue that while moral ignorance itself might not be exculpatory, epistemically justified moral ignorance might well be and that the version of normative internalism they have in mind only makes claims about the latter. But is epistemically justified moral ignorance actually exculpatory? And what makes moral ignorance epistemically justified? These questions were addressed in the final manuscript of this dissertation. In the literature on moral responsibility in the context of moral ignorance, different conditions of exculpatory moral ignorance have been discussed. Volitionists claim that an agent is only culpable for moral ignorance and any resulting wrongdoing if the ignorance can be traced back to a moment where the agent was aware of their obligation to better their moral epistemic state (procedural obligations) but wittingly failed to do so, for example, by not discussing their moral views with people who hold opposing views, by not investigating their gut feeling that something about their moral view might not be right, or by not seriously considering arguments brought forward against their moral view (Rosen, 2003, 2004). Proponents of Quality of Will, on the other hand, have argued that moral ignorance can only be exculpatory if it is the result of low accessibility of the moral truth, for example if the available moral evidence is misleading or if the moral issue itself is particularly hard (FitzPatrick, 2008; Hartford, 2019).

In the pre-registered experiment that is described in the third manuscript, we tested whether the conditions of exculpatory moral ignorance that have been discussed in the philosophical literature are sufficient for people to ascribe less blame to a morally ignorant wrongdoer. If moral ignorance turned out to be perceived as exculpatory, this would undermine the intuitiveness of some arguments raised against normative internalism. If, however, even epistemically justified moral ignorance would not be perceived as exculpatory, this would put pressure on some readings of normative internalism and consequently on respective readings of moral uncertaintism. It would also show a fundamental discrepancy between lay views and the motivation behind moral uncertaintism.

To test different claims regarding exculpatory moral ignorance, we used a modified version of a blame-updating paradigm (Monroe & Malle, 2019). In this paradigm,

participants are given ambiguous information about a wrongdoing and are asked for an initial blame judgment. Then, they receive further information about the action, for example about the intentions of the agent, and are given the possibility to revise their judgment based on the new information. Exculpation is then operationalized as the reduction in blame ascriptions between the initial and updated judgments.

We changed several aspects of this paradigm to test claims about moral ignorance. First, we adapted the information that was given between pre- and post-measurements to describe the agent's moral ignorance. This information followed a 2 (fulfilled vs. violated procedural obligations) × 2 (biased vs. unbiased moral evidence) design that closely followed volitionist and Quality of Will arguments. Second, we included a control trial to differentiate between retest effects and actual effects of moral ignorance. Last, instead of only relying on manifest change scores, we also used structural equation modeling for a latent analysis of blame updating, addressing potential concerns regarding the reliability of manifest change scores, as they have been used in previous analyses of data from blame-updating paradigms (Kievit et al., 2018; Monroe & Malle, 2019; Steyer et al., 1997).

We collected data from N = 251 participants who each underwent four moral ignorance trials and one control trial in randomized order. In the beginning of the study, participants were given a list of 39 actions and were asked to choose five that they morally reject. The selected actions were rated in terms of their moral difficulty and were then randomly assigned to the five trials. In each trial, participants gave pre- and post-ratings of a singleitem measure of act wrongness and a three-item measure of act blameworthiness. In the post-ratings of moral ignorance trials, participants were also asked to respond to a threeitem measure of belief blameworthiness, capturing how blameworthy the agent was for

the described moral ignorance.



Figure 3.1. Pre-post changes in act blameworthiness by experimental condition. Possible values ranged from -100 to 100. Positive values indicate increased blame, negative values indicate a blame reduction. Bars indicate estimated marginal means for each condition including a 95% confidence interval. Dunnett-corrected *p*-values and Cohen's *d* for each comparison against the control condition are given. As a reference, the outmost right bar indicates the blame reduction Monroe & Malle (2019) found in their first study for acting unintentionally. Reproduced from Theisen and Andow (2024).

We found that moral ignorance did not have an overall exculpatory effect. There was a net reduction in blame ascriptions for the interaction of fulfilled procedural obligations and biased moral evidence, but this reduction was not significantly different from the reduction that was observed in the control trials (Figure 3.1). The strongest predictors of exculpation were low belief blameworthiness, i.e., the agent being blameless for their

moral view, and high moral difficulty. Our findings suggested that blameless moral ignorance might in principle be perceived as exculpatory, but that the conditions that were tested in this study were not sufficient to make the described agents blameless for their moral view.

Overall, our findings lent suggestive support to critics of normative internalism since the agent's moral epistemic state did not have a relevant effect on participants' ascriptions of blame. Yet, the strongly negative association between belief blameworthiness and blame mitigation suggested that blameless moral ignorance may in principle be seen as exculpatory, but that this condition might be seldomly met in daily life. While our findings did not suggest that lay people share the assumption of normative internalism, it could therefore not be completely ruled out that moral epistemic states at least sometimes affect subjective rightness according to lay people.

4 Discussion

The question of how to make decisions under normative uncertainty, which has recently sparked debate in analytic philosophy under the label of moral uncertaintism, concerns philosophers and lay people alike. In this dissertation, I laid the groundwork for the empirical study of such choices by testing the consistencies and discrepancies between philosophical and lay concepts of normative uncertainty. Two major insights emerged from this work: First, normative uncertainty is a psychological phenomenon that is distinct from uncertainty about empirical facts, resulting in a need for a descriptive theory of choice under normative uncertainty. Second, in central aspects, the moral uncertaintist discussion of normative uncertainty does not align with lay intuitions, which must be considered in the empirical study of such uncertainty. Discrepancies between lay and moral uncertaintist concepts were identified regarding the metaethical interpretation of normative uncertainty, which affects any empirical study of normative uncertainty, and regarding the association between moral epistemic states and subjective rightness, which affects studies on how lay people think one should navigate choice under normative uncertainty

MS	Assumption	Main finding	Implication
1	Empirical-normative distinction	Empirical and normative uncertainty are separate psychological states	Need for descriptive theory of choice under normative uncertainty
2	Realism	Lay people mostly reject realist- leaning interpretations of normative uncertainty	Need for metaethically agnostic terms to describe normative uncertainty in study material
3	Internalism	Moral ignorance is not perceived as exculpatory	Doubt regarding the usefulness of studies addressing how lay people judge strategies for choice under normative uncertainty

Table 4.1. Overview of This Dissertation's Findings

I will discuss the implications of this dissertation's findings for the development of descriptive theories of choice under normative uncertainty (section 4.1), before reaching a conclusion on how moral psychology may benefit from the moral uncertaintism discourse (section 4.2).

4.1 Implications for Descriptive Theories of Choice Under Normative Uncertainty

The first assumption of moral uncertaintism was confirmed: There is moral uncertainty due to normative questions. This was shown for everyday moral cases reported by the study participants themselves and for researcher-generated moral dilemmas as they have typically been used in moral psychology. The findings reported in the first manuscript thus highlight the importance of developing a descriptive theory of choice under normative uncertainty. This theory might resemble strategies that have been discussed as part of moral uncertaintism, such as acting according to the moral view one is most certain of (Gracely, 1996; Gustafsson & Torpman, 2014) or maximizing expected moral value (MacAskill & Ord, 2020). It is also possible that lay people's choices under normative uncertainty follow principles that have so far not been considered in the moral uncertaintism literature and that need to be spelled out by moral psychology.¹

Findings from the first manuscript also suggested that, for lay people, normative uncertainty primarily takes the form of moral conflict and is less about uncertainty between discrete normative claims. While this highlights the relevance of normative uncertainty for the empirical study of moral dilemmas, it might also prompt future

¹ Consider how the empirical study of choice under uncertainty in non-moral contexts yielded models that differed from traditional decision theory, such as the seminal prospect theory (Kahneman & Tversky, 1979). A similar dynamic might unfold when developing descriptive theories of choice under normative uncertainty.

research testing whether normative uncertainty about discrete normative claims is also part of lay people's moral cognition.

While the first manuscript established that normative uncertainty can be a relevant subject of moral psychology, the remaining two manuscripts casted doubt on (1) the suitability of moral uncertaintist models as a framework to study lay people's moral thought and practice and (2) the relevance of normative internalism for lay people. Findings from the second manuscript revealed that the moral uncertaintism discourse is closely associated with moral realism, while lay people largely reject realist interpretations of normative uncertainty. Primarily, this has consequences for the language and concepts that are used to describe normative uncertainty in psychological paradigms, such as speaking of moral facts or describing normative uncertainty in percentages. Even though some moral uncertaintists have stated that their usage of realist-leaning terms does not necessarily imply that realism is a prerequisite of moral uncertaintism (e.g., Carr, 2020; Hedden, 2016; Moller, 2011; Rosenthal, 2019), using such terms in the empirical study of choice under normative uncertainty might lead to misunderstandings and bias on the part of the participants since these terms do not represent how most lay people conceive of normative uncertainty. Researchers may thus be advised to use metaethically neutral terms where possible.

Going beyond the terms that would be used in empirical studies of choice under normative uncertainty, it also needs to be discussed which role moral realism plays for the moral uncertaintist idea itself. Treating normative uncertainty analogously to empirical uncertainty might seem more plausible when assuming a fundamental similarity between empirical and normative facts (Tarsney, 2021; E. G. Williams, 2015). This would suggest that lay people's rejection of moral realism might not only conflict with the terms that moral uncertaintists have used but also with the normative internalist core of moral uncertaintism itself.

Lay people's views on normative internalism were further explored in the third manuscript. Here, it could be shown that moral epistemic states, i.e., the moral beliefs of a wrongdoer, did not affect as how blameworthy the wrongdoer was perceived. This is relevant since blameworthiness and subjective rightness have often been discussed as closely linked, so our findings might be interpreted as showing that people do not think that what an agent subjectively ought to do depends on their moral epistemic state (Harman, 2015, 2022). This suggests that lay people may not hold normative internalist intuitions.

However, not all is lost for lay normative internalism. First, the third manuscript focused on moral ignorance, which has been discussed as the limiting case of moral uncertainty (Harman, 2015, 2022). The rationale behind this was that, if normative internalism is true, moral ignorance should give the agent stronger subjective reason to commit a wrongdoing than normative uncertainty – if someone is 100% certain that physical punishment should be a central part of education, this would give them stronger subjective reason to conduct such punishment than if they were only 80% certain. If even those stronger reasons are not sufficient to affect what an agent subjectively ought to do, the same can be expected from weaker reasons arising from normative uncertainty, so the argument. Yet, as has been argued in the manuscript, it might be that the processes underlying perceptions of morally ignorant wrongdoers differ from those underlying the perception of normatively uncertain wrongdoers. While moral ignorance might be closely associated with negative perceptions of the agent's character, thereby affecting ratings of blameworthiness, the same may be less applicable to normative uncertainty (see also Biebel, 2023; Schwartz & Inbar, 2023). Thus, people might be more accepting of a wrongdoer acting according to their normative uncertainty than their moral ignorance. This would leave some room for the kind of lay normative internalism that is required for lay moral uncertaintism.

Second, the findings from the third manuscript focused on ascriptions of blame to assess whether people think that moral epistemic states affect subjective rightness. However, the link between blameworthiness and subjective rightness has been contested (Geyer, 2018; Sepielli, 2017). If subjective rightness is understood in rational rather than moral terms, an agent might do what they subjectively ought to do and still be blameworthy. Findings regarding the exculpatory potential of moral ignorance are relevant to discussions of lay normative internalism only to the extent to which subjective rightness rules out blameworthiness.

Last, even though an exculpatory effect of moral ignorance was not found, regression analysis suggested that such an effect might be expected for fully non-culpable moral ignorance, assuming a linear relationship between belief blameworthiness and act blameworthiness. Even though the conditions in our study were not sufficient for participants to perceive the described moral ignorance as fully non-culpable, it cannot be ruled out that such cases exist; again, leaving some room for lay normative internalism.

In future studies, lay people's intuitions regarding different versions of normative internalism need to be tested both using descriptions of moral ignorance as well as descriptions of normative uncertainty and by considering different hypotheses regarding the association between subjective rightness and blameworthiness. The current findings suggest that lay people might not share the normative internalist core of moral uncertaintism.

The potential discrepancy between normative internalism, on which the moral uncertaintism project builds, and lay people's normative externalist intuitions does not necessarily affect studies on how lay people make choices under normative uncertainty. Lay people still have to make such choices, which can then be the subject of a descriptive theory of choice under normative uncertainty, even if they do not think that there are better and worse ways to make such decisions. Yet, people's rejection of normative internalism would affect studies focusing on how lay people judge different strategies to make choices under normative uncertainty. Do lay people think that one should refrain from certain actions just because one is normatively uncertain? Do they think that difficult moral cases require moral compromise out of moral uncertaintist considerations? The study of such questions becomes less relevant if lay people's rejection of normative internalism is further confirmed.

4.2 Conclusion

At the beginning of this dissertation, it was asked whether the search for descriptive theories of choice under normative uncertainty can benefit from the work that has already been done on normative theories about such choices. The present findings have shown the limitations of this approach while simultaneously confirming the need for such descriptive theories.

A central aspect of many moral uncertaintist models is to consider certainty separately from considerations of the respective moral stakes (i.e., conditional moral value; MacAskill et al., 2020; see also M. Smith, 2002). The previous discussion of normative uncertainty in the context of process-dissociation models and the moral trade-off model has revealed how considering this separation might allow for more specific inferences about participant's moral cognition, for example, by allowing to differentiate between different motivations behind moral compromise. Even though findings from the second manuscript urge caution not to use realist-leaning terms for such measurements, it might be possible to assess people's certainty regarding the relative importance of the competing moral arguments in moral conflict cases without irritating anti-realist-leaning participants. If this succeeds, the separate consideration of certainty and conditional moral value might be one way in which moral uncertaintism can advance the study of choice under normative uncertainty.

A second way in which moral psychology might benefit from moral uncertaintism is via new research questions that are inspired by the moral uncertaintism literature. It has already been mentioned that researchers may focus on which strategy regarding choice under normative uncertainty is seen as better or worse by lay people; however, lay people's potential rejection of normative internalism undermines the relevance of such questions. Yet, there are further hypotheses that can be derived from the moral uncertaintism discourse that may enhance our understanding of moral cognition and metacognition. For example, researchers may refer to the motivational aspects of moral uncertaintism: Critics have argued that moral uncertaintism requires agents to be motivated to do what is morally right, regardless of what it is that is morally right. This has been criticized as *fetishistic*. These critics argue that agents should be motivated by the right-making features of an action – such as that it benefits others, reduces harm, or is an expression of respect, etc. – and not by the rightness of the action itself (Hedden, 2016; M. Smith, 1994; Weatherson, 2014, 2019; but see among others Carbonell, 2013; Rosenthal, 2019; Sepielli, 2016). From a psychological perspective, it could be studied to what degree lay people are motivated by rightness in itself and how they evaluate moral fetishists. This would further highlight how lay people's views on morality align with the moral-theoretic underpinnings of moral uncertaintism.

At the very least, it can be stated that the recent surge in moral uncertaintist literature has already inspired first researchers to consider normative uncertainty in their empirical work (Costa-Gomes & Schoenegger, 2023; Dietrich et al., 2019; Jabarian, 2020). This dissertation is a further of those examples. The present findings suggest that, despite the discrepancies between philosophical and lay concepts regarding normative uncertainty, normative uncertainty is a relevant subject of moral psychology that deserves further attention.

References

- Alsaad, A., Elrehail, H., & Saif-Alyousfi, A. Y. H. (2022). The interaction among religiosity, moral intensity and moral certainty in predicting ethical consumption: A study of Muslim consumers. *International Journal of Consumer Studies*, 46(2), 406–418. https://doi.org/10.1111/ijcs.12688
- Altham, J. E. T. (1983-1984). Ethics of risk. *Proceedings of the Aristotelian Society, New Series*, *84*, 15–29.
- American Psychological Association. (2020). *Publication manual of the American Psychological Association: The official guide to APA style* (7th ed.). American Psychological Association. https://doi.org/10.1037/0000165-000
- Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J.-F., & Rahwan, I. (2018). The Moral Machine experiment. *Nature*, 563(7729), 59–64. https://doi.org/10.1038/s41586-018-0637-6
- Bago, B., & De Neys, W. (2019). The intuitive greater good: Testing the corrective dual process model of moral cognition. *Journal of Experimental Psychology. General*, 148(10), 1782–1801. https://doi.org/10.1037/xge0000533
- Bai, H., Euh, H., Federico, C. M., & Borgida, E. (2021). Thou shalt not kill, unless it is not a human: Target dehumanization may influence decision difficulty and response patterns for moral dilemmas. *Social Cognition*, 39(6), 657–686. https://doi.org/10.1521/soco.2021.39.6.657
- Behnke, A., Strobel, A., & Armbruster, D. (2020). When the killing has been done: Exploring associations of personality with third-party judgment and punishment of homicides in moral dilemma scenarios. *PloS One*, 15(6), e0235253. https://doi.org/10.1371/journal.pone.0235253
- Białek, M., & De Neys, W. (2016). Conflict detection during moral decision-making: evidence for deontic reasoners' utilitarian sensitivity. *Journal of Cognitive Psychology*, 28(5), 631–639. https://doi.org/10.1080/20445911.2016.1156118
- Białek, M., & De Neys, W. (2017). Dual processes and moral conflict: Evidence for deontological reasoners' intuitive utilitarian sensitivity. *Judgment and Decision Making*, 12(2), 148–167.
- Biebel, N. (2023). Reasonable standards and exculpating moral ignorance. *Philosophical Studies*. Advance online publication. https://doi.org/10.1007/s11098-023-02054-x
- Bogosian, K. (2017). Implementation of moral uncertainty in intelligent machines. *Minds* and Machines, 27(4), 591–608. https://doi.org/10.1007/s11023-017-9448-z
- Bourget, D., & Chalmers, D. J. (2023). Philosophers on philosophy: The 2020 PhilPapers survey. *Philosophers' Imprint*, 23(1). https://doi.org/10.3998/phimp.2109

- Brannon, S. M., DeJong, A., & Gawronski, B. (2019). Determinants of lateral attitude change: The roles of object relatedness, attitude certainty, and moral conviction. *Social Cognition*, 37(6), 624–658. https://doi.org/10.1521/soco.2019.37.6.624
- Brennan, S. (1995). Thresholds for rights. *The Southern Journal of Philosophy*, *33*(2), 143–168. https://doi.org/10.1111/j.2041-6962.1995.tb00737.x
- Brink, D. O. (1989). *Moral realism and the foundation of ethics. Cambridge studies in philosophy*. Cambridge University Press.
- Bykvist, K. (2017). Moral uncertainty. *Philosophy Compass*, 12(3), e12408. https://doi.org/10.1111/phc3.12408
- Carbonell, V. (2013). De dicto desires and morality as fetish. *Philosophical Studies*, *163*(2), 459–477.
- Carmona-Perera, M., Caracuel, A., Pérez-García, M., & Verdejo-García, A. (2015). Brief Moral Decision-Making Questionnaire: A Rasch-derived short form of the Greene dilemmas. *Psychological Assessment*, 27(2), 424–432. https://doi.org/10.1037/pas0000049
- Carr, J. R. (2015). Subjective ought. *Ergo*, 2(20201214). https://doi.org/10.3998/ergo.12405314.0002.027
- Carr, J. R. (2020). Normative uncertainty without theories. *Australasian Journal of Philosophy*, 98(4), 747–762. https://doi.org/10.1080/00048402.2019.1697710
- Chandler, J. (2017). Descriptive decision theory. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/win2017/entries/decision-theory-descriptive/
- Chen, K., Zhang, H., Friedman, M., & Schlegel, R. J. (2022). The authentic catch-22: Following the true self promotes decision satisfaction in moral dilemmas. *Journal* of Experimental Social Psychology, 102, 104376. https://doi.org/10.1016/j.jesp.2022.104376
- Conway, P., & Gawronski, B. (2013). Deontological and utilitarian inclinations in moral decision making: A process dissociation approach. *Journal of Personality and Social Psychology*, 104(2), 216–235. https://doi.org/10.1037/a0031021
- Costa-Gomes, M., & Schoenegger, P. (2023). Charitable giving under normative uncertainty: Experimental evidence on the behavioral impact of normative expert advice. SSRN Electronic Journal. Advance online publication. https://doi.org/10.2139/ssrn.4389737
- Côté, N. (2023). There is no such thing as expected moral choice-worthiness. *Canadian Journal of Philosophy*, 53(1), 1–20. https://doi.org/10.1017/can.2023.18
- Crutchfield, P., Scheall, S., Rzeszutek, M. J., Brown, H. D., & Cardoso Sao Mateus, C. (2023). Ignorance and moral judgment: Testing the logical priority of the epistemic. *Consciousness and Cognition*, 108, 103472. https://doi.org/10.1016/j.concog.2023.103472
- DeGrazia, D. (2008). Moral status as a matter of degree? *The Southern Journal of Philosophy*, *46*, 181–198.

- Dietrich, F., Fleurbaey, M., & Jabarian, B. (2019, July 1). *The normative uncertainty survey*. 2nd Workshop on Global Priorities Research. Oxford University, Oxford.
- Domian, J. (2008, October 29). Domian: Moralische Konflikte. 1 Live.
- Doyle, A., & Ericson, D. (2003). *Risk and morality*. University of Toronto Press. https://doi.org/10.3138/9781442679382
- Ellemers, N., van der Toorn, J., Paunov, Y., & van Leeuwen, T. (2019). The psychology of morality: A review and analysis of empirical studies published from 1940 through 2017. *Personality and Social Psychology Review*, 23(4), 332–366. https://doi.org/10.1177/10888683188117
- Eriksson, J., & Francén Olinder, R. (2016). Non-cognitivism and the classification account of moral uncertainty. *Australasian Journal of Philosophy*, 94(4), 719– 735. https://doi.org/10.1080/00048402.2016.1149199
- Evans, J. S. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255–278. https://doi.org/10.1146/annurev.psych.59.103006.093629
- Evans, J. S. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, 8(3), 223–241. https://doi.org/10.1177/1745691612460685
- FitzPatrick, W. J. (2008). Moral responsibility and normative ignorance: Answering a new skeptical challenge. *Ethics*, 118(4), 589–613. https://doi.org/10.1086/589532
- Fleischhut, N., Meder, B., & Gigerenzer, G. (2017). Moral hindsight. *Experimental Psychology*, *64*(2), 110–123. https://doi.org/10.1027/1618-3169/a000353
- Foot, P. (1967). The problem of abortion and the doctrine of the double effect. *Oxford Review*, *5*, 5–15.
- Gawronski, B., Armstrong, J., Conway, P., Friesdorf, R., & Hütter, M. (2017). Consequences, norms, and generalized inaction in moral dilemmas: The CNI model of moral decision-making. *Journal of Personality and Social Psychology*, *113*(3), 343–376. https://doi.org/10.1037/pspa0000086
- Gerpott, F. H., Balliet, D., Columbus, S., Molho, C., & Vries, R. E. de (2018). How do people think about interdependence? A multidimensional model of subjective outcome interdependence. *Journal of Personality and Social Psychology*, *115*(4), 716–742. https://doi.org/10.1037/pspp0000166
- Geyer, J. (2018). Moral uncertainty and moral culpability. *Utilitas*, *30*(4), 399–416. https://doi.org/10.1017/S0953820818000080
- Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. Cognition, 106(3), 1339–1366. https://doi.org/10.1016/j.cognition.2007.06.007
- Goodwin, G. P., & Darley, J. M. (2012). Why are some moral beliefs perceived to be more objective than others? *Journal of Experimental Social Psychology*, 48(1), 250–256. https://doi.org/10.1016/j.jesp.2011.08.006

- Gracely, E. J. (1996). On the noncomparability of judgments made by different ethical theories. *Metaphilosophy*, *27*(3), 327–332. https://doi.org/10.1111/j.1467-9973.1996.tb00212.x
- Greene, J. D. (2023). The dual-process theory of moral judgment does not deny that people can make compromise judgments. *Proceedings of the National Academy of Sciences of the United States of America*, 120(6), e2220396120. https://doi.org/10.1073/pnas.2220396120
- Greene, J. D., & Haidt, J. (2002). How (and where) does moral judgment work? *Trends in Cognitive Sciences*, 6(12), 517–523.
- Greene, J. D., Nystrom, L. E., Engell, A. D., Darley, J. M., & Cohen, J. D. (2004). The neural bases of cognitive conflict and control in moral judgment. *Neuron*, 44(2), 389–400. https://doi.org/10.1016/j.neuron.2004.09.027
- Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, 293(5537), 2105–2108. https://doi.org/10.1126/science.1062191
- Greenhalgh, T. (2021). Moral uncertainty: A case study of Covid-19. *Patient Education* and Counseling, 104(11), 2643–2647. https://doi.org/10.1016/j.pec.2021.07.022
- Guerrero, A. A. (2007). Don't know, don't kill: Moral Ignorance, culpability, and caution. *Philosophical Studies*, *136*(1), 59–97. https://doi.org/10.1007/s11098-007-9143-7
- Gustafsson, J. E. (2022). Second thoughts about My Favourite Theory. *Pacific Philosophical Quarterly*, Article papq.12408. Advance online publication. https://doi.org/10.1111/papq.12408
- Gustafsson, J. E., & Torpman, O. (2014). In defence of My Favourite Theory. *Pacific Philosophical Quarterly*, 95(2), 159–174. https://doi.org/10.1111/papq.12022
- Guzmán, R. A., Barbato, M. T., Sznycer, D., & Cosmides, L. (2022). A moral trade-off system produces intuitive judgments that are rational and coherent and strike a balance between conflicting moral values. *Proceedings of the National Academy* of Sciences of the United States of America, 119(42), e2214005119. https://doi.org/10.1073/pnas.2214005119
- Hanselmann, M., & Tanner, C. (2008). Taboos and conflicts in decision making: Sacred values, decision difficulty, and emotions. *Judgment and Decision Making*, 3(1), 51–63.
- Hansson, S. O. (2013). *The ethics of risk*. Palgrave Macmillan UK. https://doi.org/10.1057/9781137333650
- Harman, E. (2015). The irrelevance of moral uncertainty. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 10, pp. 53–79). Oxford University Press.
- Harman, E. (2022). Ethics is hard! What follows? In D. K. Nelkin & D. Pereboom (Eds.), *The Oxford handbook of moral responsibility* (pp. 327–347). Oxford University Press.

- Hartford, A. (2019). Moral and factual ignorance: A quality of will parity. *Ethical Theory and Moral Practice*, *22*(5), 1087–1102. https://doi.org/10.1007/s10677-019-10043-5
- Hayenhjelm, M., & Wolff, J. (2012). The moral problem of risk impositions: A survey of the literature. *European Journal of Philosophy*, 20(S1). https://doi.org/10.1111/j.1468-0378.2011.00482.x
- Hedden, B. (2016). Does MITE make right? On decision-making under normative uncertainty. In R. Shafer-Landau (Ed.), Oxford studies in metaethics. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198784647.001.0001
- Hicks, A. (2018). Moral uncertainty and value comparison. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (pp. 161–183). Oxford University Press.
- Hofmann, W., Wisneski, D. C., Brandt, M. J., & Skitka, L. J. (2014). Morality in everyday life. *Science*, *345*(6202), 1340–1343. https://doi.org/10.1126/science.1251560
- Jabarian, B. (2020, May 30). *Moral uncertainty, social values and economic preferences* [Paper presentation]. Princeton University, Princeton, NJ, USA.
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, 30(5), 513–541. https://doi.org/10.1016/0749-596X(91)90025-F
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263. https://doi.org/10.2307/1914185
- Kievit, R. A., Brandmaier, A. M., Ziegler, G., van Harmelen, A.-L., Mooij, S. M. M. de, Moutoussis, M., Goodyer, I. M., Bullmore, E., Jones, P. B., Fonagy, P., Lindenberger, U., & Dolan, R. J. (2018). Developmental cognitive neuroscience using latent change score models: A tutorial and applications. *Developmental Cognitive Neuroscience*, 33, 99–117. https://doi.org/10.1016/j.dcn.2017.11.007
- Kirfel, L., & Hannikainen, I. R. (2023). Why blame the ostrich? Understanding culpability for willful ignorance. In S. Magen & K. Prochownik (Eds.), Advances in experimental philosophy of law (pp. 75–98). Bloomsbury Academic.
- Kissinger-Knox, A., Aragon, P., & Mizrahi, M. (2018). Does non-moral ignorance exculpate? Situational awareness and attributions of blame and forgiveness. *Acta Analytica*, 33, 161–179.
- Kneer, M. (2018). Perspective and epistemic state ascriptions. *Review of Philosophy and Psychology*, 9(2), 313–341. https://doi.org/10.1007/s13164-017-0361-4
- Körner, A., Deutsch, R., & Gawronski, B. (2020). Using the CNI model to investigate individual differences in moral dilemma judgments. *Personality & Social Psychology Bulletin*, 46(9), 1392–1407. https://doi.org/10.1177/0146167220907203
- Kortenkamp, K. V., & Moore, C. F. (2014). Ethics under uncertainty: The morality and appropriateness of utilitarianism when outcomes are uncertain. *American Journal* of Psychology, 127(3), 367. https://doi.org/10.5406/amerjpsyc.127.3.0367

- Kuhnert, B., Lindner, F., Bentzen, M. M., & Ragni, M. (2017). Perceived difficulty of moral dilemmas depends on their causal structure: A formal model and preliminary results. 39th Annual Meeting of the Cognitive Science Society.
- Lee, M., Sul, S., & Kim, H. (2018). Social observation increases deontological judgments in moral dilemmas. *Evolution and Human Behavior*, 39(6), 611–621. https://doi.org/10.1016/j.evolhumbehav.2018.06.004
- Lewens, T. (Ed.). (2007). Risk: Philosophical perspectives. Routledge.
- Lockhart, T. (2000). Moral uncertainty and its consequences. Oxford University Press.
- MacAskill, W. (2014). Normative uncertainty [Thesis]. University of Oxford, Oxford.
- MacAskill, W. (2016). Normative uncertainty as a voting problem. *Mind*, 125(500), 967–1004. https://doi.org/10.1093/mind/fzv169
- MacAskill, W., Bykvist, K., & Ord, T. (Eds.). (2020). *Moral uncertainty*. Oxford University Press.
- MacAskill, W., & Ord, T. (2020). Why maximize expected choice-worthiness? *Noûs*, 54(2), 327–353. https://doi.org/10.1111/nous.12264
- Makins, N. (2021). Attitudinal ambivalence: Moral uncertainty for non-cognitivists. *Australasian Journal of Philosophy*, 1–15. https://doi.org/10.1080/00048402.2021.1908380
- Mandel, D. R., & Vartanian, O. (2008). Taboo or tragic: effect of tradeoff type on moral choice, conflict, and confidence. *Mind & Society*, 7(2), 215–226. https://doi.org/10.1007/s11299-007-0037-3
- Martinho, A., Kroesen, M., & Chorus, C. (2021). Computer says I don't know: An empirical approach to capture moral uncertainty in artificial intelligence. *Minds and Machines*, *31*(2), 215–237. https://doi.org/10.1007/s11023-021-09556-9
- Mason, H. E. (Ed.). (1996). Moral dilemmas and moral theory. Oxford University Press.
- Mata, A. (2019). Social metacognition in moral judgment: Decisional conflict promotes perspective taking. *Journal of Personality and Social Psychology*, 117(6), 1061– 1082. https://doi.org/10.1037/pspa0000170
- McConnell, T. (2022). Moral dilemmas. In Zalta, Edward N. & Nodelman, Uri (Ed.), *The Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/fall2022/entries/moral-dilemmas/
- Meder, B., Fleischhut, N., Krumnau, N.-C., & Waldmann, M. R. (2019). How should autonomous cars drive? A preference for defaults in moral judgments under risk and uncertainty. *Risk Analysis*, 39(2), 295–314. https://doi.org/10.1111/risa.13178
- Merlhiot, G., Mermillod, M., Le Pennec, J.-L., Dutheil, F., & Mondillon, L. (2018). Influence of uncertainty on framed decision-making with moral dilemma. *PloS One*, 13(5), e0197923. https://doi.org/10.1371/journal.pone.0197923
- Moller, D. (2011). Abortion and moral risk. *Philosophy*, *86*(3), 425–443. https://doi.org/10.1017/S0031819111000222

- Monroe, A. E., & Malle, B. F. (2019). People systematically update moral judgments of blame. *Journal of Personality and Social Psychology*, 116(2), 215–236. https://doi.org/10.1037/pspa0000137
- Moore, M. S. (2010). *Placing blame*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199599493.001.0001
- New York Times Magazine. (1999). *The Ethicist*. https://www.nytimes.com/column/theethicist
- Newberry, T., & Ord, T. (2021). *The parliamentary approach to moral uncertainty*. Technical Report #2021-2. Future of Humanity Institute, University of Oxford.
- Ng, N. L., Luke, D. M., & Gawronski, B. (2023). Moral judgment under uncertainty: A CNI model analysis. *European Journal of Social Psychology*, *53*(6), 1055–1077. https://doi.org/10.1002/ejsp.2952
- Nissan-Rozen, I. (2015). Against moral hedging. *Economics and Philosophy*, 31(3), 349–369. https://doi.org/10.1017/S0266267115000206
- O'Connor, A. M. (1995). Validation of a Decisional Conflict Scale. *Medical Decision Making*, 15, 25–30.
- Oddie, G. (1994). Moral uncertainty and human embryo experimentation. In K. W. M. Fulford, G. R. Gillett, & J. M. Soskice (Eds.), *Medicine and moral reasoning* (pp. 144–161). Cambridge University Press.
- Parfit, D. (1984). Reasons and persons. Oxford University Press.
- Parfit, D. (2011). On what matters. Oxford University Press.
- Petrocelli, J. V., Tormala, Z. L., & Rucker, D. D. (2007). Unpacking attitude certainty: Attitude clarity and attitude correctness. *Journal of Personality and Social Psychology*, 92(1), 30–41. https://doi.org/10.1037/0022-3514.92.1.30
- Podgorski, A. (2020). Normative uncertainty and the dependence problem. *Mind*, *129*(513), 43–70. https://doi.org/10.1093/mind/fzz048
- Poliquin, C. W. (2010). Protected values, range effects, guilt, and tradeoff difficulty in moral decision making. *Penn Journal of Philosophy, Politics & Economics*, 5(1), Article 5. https://repository.upenn.edu/spice/vol5/iss1/5
- Pölzler, T., & Wright, J. C. (2020). Anti-realist pluralism: A new approach to folk metaethics. *Review of Philosophy and Psychology*, 11(1), 53–82. https://doi.org/10.1007/s13164-019-00447-8
- Railton, P. (1986). Moral realism. *The Philosophical Review*, 95(2), 163–207. https://doi.org/10.2307/2185589
- Robichaud, P., & Wieland, J. W. (Eds.). (2017). *Responsibility: The epistemic condition*. Oxford University Press.
- Robinson, P. (2022). Is normative uncertainty irrelevant if your descriptive uncertainty depends on it? *Pacific Philosophical Quarterly*, *103*(4), 874–899. https://doi.org/10.1111/papq.12379

- Robinson, P., & Steele, K. (2022). Moral uncertainty, noncognitivism, and the multiobjective story. *Noûs*, Article nous.12435. Advance online publication. https://doi.org/10.1111/nous.12435
- Rosas, A., Bermúdez, J. P., & Aguilar-Pardo, D. (2019). Decision conflict drives reaction times and utilitarian responses in sacrificial dilemmas. *Judgment and Decision Making*, 14(5), 555–564. https://doi.org/10.1017/S193029750000485X
- Rosen, G. (2003). Culpability and ignorance. *Proceedings of the Aristotelian Society*, 103(1), 61–84.
- Rosen, G. (2004). Skepticism about moral responsibility. *Philosophical Perspectives*, *18*(Ethics), 295–313.
- Rosenthal, C. (2019). *Ethics for fallible people* [Dissertation]. New York University, New York, NY.
- Rosenthal, C. (2021). What decision theory can't tell us about moral uncertainty. *Philosophical Studies*, *178*(10), 3085–3105. https://doi.org/10.1007/s11098-020-01571-3
- Ross, J. (2006). Rejecting ethical deflationism. *Ethics*, *116*(4), 742–768. https://doi.org/10.1086/505234
- Ross, W. D. (1930, 2002). *The right and the good: Some problems in ethics*. Oxford University Press.
- Ryazanov, A. A., Knutzen, J., Rickless, S. C., Christenfeld, N. J. S., & Nelkin, D. K. (2018). Intuitive probabilities and the limitation of moral imagination. *Cognitive Science*, 42 Suppl 1, 38–68. https://doi.org/10.1111/cogs.12598
- Ryazanov, A. A., Wang, S. T., Rickless, S. C., McKenzie, C. R. M., & Nelkin, D. K. (2021). Sensitivity to shifts in probability of harm and benefit in moral dilemmas. *Cognition*, 209, 104548. https://doi.org/10.1016/j.cognition.2020.104548
- Savage, L. J. (1954). The foundations of statistics. John Wiley & Sons.
- Schwartz, S. A., & Inbar, Y. (2023). Is it good to feel bad about littering? Conflict between moral beliefs and behaviors for everyday transgressions. *Cognition*, 236, 105437. https://doi.org/10.1016/j.cognition.2023.105437
- Sepielli, A. (2009). What to do when you don't know what to do. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (pp. 5–28). Oxford University Press.
- Sepielli, A. (2010). 'Along an imperfectly-lighted path': Practical rationality and normative uncertainty [Dissertation]. The State University of New Jersey, New Brunswick, NJ.
- Sepielli, A. (2016). Moral uncertainty and fetishistic motivation. *Philosophical Studies*, 173(11), 2951–2968. https://doi.org/10.1007/s11098-016-0645-z
- Sepielli, A. (2017). How moral uncertaintism can be both true and interesting. In M. C. Timmons (Ed.), Oxford studies in normative ethics (pp. 98–116). Oxford University Press. https://doi.org/10.1093/oso/9780198808930.003.0006
- Shaddy, F., Fishbach, A., & Simonson, I. (2021). Trade-offs in choice. *Annual Review of Psychology*, 72, 181–206. https://doi.org/10.1146/annurev-psych-072420-125709

Shafer-Landau, R. (2003). Moral realism: A defence. Clarendon.

- Shamoun, S., & Svenson, O. (2002). Value conflict and post-decision consolidation. *Scandinavian Journal of Psychology*, *43*, 325–333.
- Shenhav, A., & Greene, J. D. (2010). Moral judgments recruit domain-general valuation mechanisms to integrate representations of probability and magnitude. *Neuron*, 67(4), 667–677. https://doi.org/10.1016/j.neuron.2010.07.020
- Shivnekar, R., & Srivastava, N. (2023). Measuring moral vacillation. *Proceedings of the Annual Conference of the Cognitive Science Society*, 45, 464–470.
- Shou, Y., Olney, J., Smithson, M., & Song, F. (2020). Impact of uncertainty and ambiguous outcome phrasing on moral decision-making. *PloS One*, 15(5), e0233127. https://doi.org/10.1371/journal.pone.0233127
- Shou, Y., & Song, F. (2017). Decisions in moral dilemmas: The influence of subjective beliefs in outcome probabilities. *Judgment and Decision Making*, *12*(5), 481–490.
- Sinnott-Armstrong, W. (1988). Moral dilemmas. Blackwell.
- Skovgaard-Olsen, N., & Klauer, K. C. (2024). Norm conflicts and morality: The CNIS conflict model of moral decision-making [Unpublished manuscript].
- Smith, H. M. (2010). Subjective rightness. *Social Philosophy and Policy*, *27*(2), 64–110. https://doi.org/10.1017/S0265052509990161
- Smith, M. (1994). The moral problem. Blackwell.
- Smith, M. (2002). Evaluation, uncertainty and motivation. *Ethical Theory and Moral Practice*, 5(3), 305–320. https://doi.org/10.1023/A:1019675327207
- Spälti, A. K., Brandt, M. J., & Zeelenberg, M. (2019). The effects of decision time on perceptions of decisions and decision makers in (moral) trade-off scenarios. Preprint. https://doi.org/10.31234/osf.io/37t2r
- Staffel, J. (2019). Expressivism, normative uncertainty, and arguments for probabilism. In T. Gendler & J. Hawthorne (Eds.), Oxford studies in epistemology (161-189). Oxford University Press.
- Steele, K., & Stefánsson, H. O. (2020). Decision theory. In E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy. https://plato.stanford.edu/archives/win2020/entries/decision-theory/
- Steyer, R., Eid, M., & Schwenkmezger, P. (1997). Modeling true intraindividual change: True change as a latent variable. *Methods of Psychological Research Online*, 2(1), 21–33.
- Süddeutsche Zeitung Magazin. (2006). *Die Gewissensfrage*. https://szmagazin.sueddeutsche.de/tag/die-gewissensfrage
- Tarsney, C. (2017). *Rationality and moral risk: A moderate defense of hedging* [Dissertation]. University of Maryland.
- Tarsney, C. (2018a). Intertheoretic value comparison: A modest proposal. Journal of Moral Philosophy, 15(3), 324–344. https://doi.org/10.1163/17455243-20170013

- Tarsney, C. (2018b). Moral uncertainty for deontologists. *Ethical Theory and Moral Practice*, 21(3), 505–520. https://doi.org/10.1007/s10677-018-9924-4
- Tarsney, C. (2019). Normative uncertainty and social choice. *Mind*, *128*(512), 1285–1308. https://doi.org/10.1093/mind/fzy051
- Tarsney, C. (2021). Normative Externalism by Brian Weatherson. *Mind*, *130*(519), 1018–1028. https://doi.org/10.1093/mind/fzaa018
- Theisen, M. (2023). Metaethical intuitions in lay concepts of normative uncertainty. *Philosophical Psychology*. Advance online publication. https://doi.org/10.1080/09515089.2023.2174422
- Theisen, M., & Andow, J. (2024). *The exculpatory potential of moral ignorance: Evidence from a blame-updating paradigm*. Manuscript submitted for publication.
- Tilley, J. J. (2000). Cultural relativism. Human Rights Quarterly, 22(2), 501-547.
- Tversky, A., & Kahnemann, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*(4481), 453–458. https://doi.org/10.1126/science.7455683
- Vega, S., Mata, A., Ferreira, M. B., & Vaz, A. R. (2020). Metacognition in moral decisions: Judgment extremity and feeling of rightness in moral intuitions. *Thinking & Reasoning*, 1–18. https://doi.org/10.1080/13546783.2020.1741448
- Wagner, J. M. J., Pölzler, T., & Wright, J. C. (2021). Implicit metaethical intuitions: Validating and employing a new IAT procedure. *Review of Philosophy and Psychology*. Advance online publication. https://doi.org/10.1007/s13164-021-00572-3
- Weatherson, B. (2014). Running risks morally. *Philosophical Studies*, *167*(1), 141–163. https://doi.org/10.1007/s11098-013-0227-2
- Weatherson, B. (2019). Normative externalism. Oxford University Press.
- Williams, B. (1981). Moral luck. Cambridge University Press.
- Williams, E. G. (2015). The possibility of an ongoing moral catastrophe. *Ethical Theory* and Moral Practice, 18(5), 971–982. https://doi.org/10.1007/s10677-015-9567-7
- Żuradzki, T. (2016). Meta-reasoning in making moral decisions under normative uncertainty. In D. Mohammed & M. Lewiński (Eds.), *Argumentation and reasoned action*. College Publications.

Appendix

- A-1 Manuscript 1: Uncertain Facts or Uncertain Values? Testing the Distinction Between Empirical and Normative Uncertainty in Moral Judgments
- A-2 Manuscript 2: Metaethical Intuitions in Lay Concepts of Normative Uncertainty
- A-3 Manuscript 3: The Exculpatory Potential of Moral Ignorance: Evidence From a Blame-Updating Paradigm

A-1 Manuscript 1 – Uncertain Facts or Uncertain Values? Testing the Distinction Between Empirical and Normative Uncertainty in Moral Judgments

Theisen, M. & Germar, M. (in press). Uncertain facts or uncertain values? Testing the distinction between empirical and normative uncertainty in moral judgments. *Cognitive Science*. https://doi.org/10.1111/cogs.13422

Uncertain Facts or Uncertain Values? Testing the Distinction Between Empirical and Normative Uncertainty in Moral Judgments

Maximilian Theisen¹, Markus Germar²

¹ Department of Psychology, Heidelberg University, Heidelberg, Germany

² Institute of Psychology, University of Hildesheim, Hildesheim, Germany

Author Note

This work was partly supported by an individual fellowship of the Baden-Württemberg *Landesgraduiertenförderung* Program from Heidelberg University, Germany. The funding source was not involved in study design, data collection, analysis or interpretation, writing of the report, or in the decision to submit the article for publication. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. The authors would like to thank the study participants for sharing their personal moral dilemmas and the expert philosophers for contributing to the item pool. The data that support the findings of this paper are openly available on OSF at https://osf.io/fjp67/.

Correspondence concerning this article should be addressed to Maximilian Theisen, Department of Psychology, Heidelberg University, Hauptstraße 47-51, D-69117 Heidelberg, Germany, email: maximilian.theisen@psychologie.uni-heidelberg.de.

Abstract

People can be uncertain in their moral judgments. Philosophers have argued that such uncertainty can either refer to the underlying empirical facts (empirical uncertainty) or to the normative evaluation of these facts itself (normative uncertainty). Psychological investigations of this distinction, however, are rare. In this paper, we combined factor-analytical and experimental approaches to show that empirical and normative uncertainty describe two related but different psychological states. In Study 1, we asked N = 265 participants to describe a case of moral uncertainty and to rate different aspects of their uncertainty about this case. Across this wide range of moral scenarios, our items loaded onto three reliable factors: Lack of information, unclear consequences, and normative uncertainty. In Study 2, we confirmed this factor structure using predefined stimulus material. N = 402 participants each rated eight scenarios that systematically varied in their degree of uncertainty regarding the consequences of the described actions and in the value conflict that was inherent to them. The empirical uncertainty factors were mainly affected by the introduction of uncertainty regarding consequences and the normative uncertainty factor was mainly affected by the introduction of value conflict. Our studies provide evidence that the distinction between empirical and normative uncertainty accurately describes a psychological reality. We discuss the relevance of our findings for research on moral judgments and decision-making, and folk metaethics.

Keywords: moral uncertainty, normative uncertainty, moral cognition, moral reasoning, moral judgment, ethics

Introduction

Would it be morally right to make vaccinations mandatory (Clarkson & Jasper, 2022; Lowth, 2020; Savulescu, 2021)? Complex moral judgments can be subject to uncertainty, meaning that people may not feel fully convinced or confident in their attitude regarding these issues (Tormala & Rucker, 2018). Such uncertainty has been called *moral uncertainty* (Alsaad, 2021; Lockhart, 2000; MacAskill et al., 2020). It is associated with more frequent rethinking as well as revising of moral judgments (Vega et al., 2020), with higher awareness of other possible moral views, with less projecting of one's own moral views onto others (Mata, 2019), and – in interaction with low attitudinal ambivalence – decreased attitude stability (Luttrell et al., 2020). Additionally, moral uncertainty might weaken the link between one's moral views and the intention to act accordingly (Alsaad, 2021), and increase the tendency to delegate the corresponding decision (Poliquin, 2010).

Despite its relevance for moral judgment and decision-making, the object of moral uncertainty remains ambiguous. What exactly are people uncertain about when they have doubt regarding the morality of, say, mandatory vaccinations? Philosophers have long distinguished between two sources of such uncertainty (Lockhart, 2000; MacAskill et al., 2020; Nissan-Rozen, 2015). First, they argue, uncertainty about the morality of a specific act might arise from ignorance of the underlying empirical facts (*empirical uncertainty*; see also Crutchfield et al., 2023; Hansson, 2013). For example, people might be uncertain regarding mandatory vaccinations because they do not know the severeness of the disease that the vaccine protects against, the safeness of the vaccine, or the consequences such a policy would have for society (Savulescu, 2021). Second, moral uncertainty might arise from unclear or conflicting normative evaluations. Even if people had full information about the circumstances and effects of a mandatory vaccination policy, they might still wonder: Do reasons of public health outweigh reasons of health autonomy? Is it generally acceptable to expose other people to a risk of side effects? Such uncertainty regarding the normative evaluation of the empirical facts rather than the facts themselves has been called *normative uncertainty* (Table 1).¹

Philosophers usually agree on the distinction between empirical and normative uncertainty. Yet, the assumption of two different forms of moral uncertainty has not been put to the test directly. In this paper, we investigate whether empirical and normative uncertainty are two distinct psychological states that refer to different objects of uncertainty, as proposed in the philosophical literature.

Uncertain Facts: Empirical Uncertainty in Moral Judgments

We distinguish between empirical uncertainty that is morally relevant as it results in moral uncertainty, and empirical uncertainty that is morally irrelevant as it does not result in moral uncertainty. Whether empirical uncertainty about a fact is morally relevant may vary interindividually, as some aspects might be morally relevant to some people but irrelevant to others (Kodapanakkal et al., 2022; Skitka, 2010). In the case of mandatory vaccinations, a textbook deontologist might judge that restricting health autonomy is wrong in itself, regardless of its consequences (Clarkson & Jasper, 2022; Lowth, 2020). Even though they might be uncertain regarding the outcomes of such a policy, this would not make them *morally* uncertain and thus, their empirical uncertainty would not be considered morally relevant. Conversely, a consequentialist would link their moral judgment of such a policy to its overall consequences. Being uncertain regarding these consequences would therefore lead to uncertainty in their moral judgment of the policy. In contrast to the deontologist, the consequentialist's empirical uncertainty would be considered morally relevant. In the following, we will use the term *empirical uncertainty* to refer to morally relevant empirical uncertainty only.

¹ We use the term *normative* not in the narrow sense of referring to norms but in a broader sense, referring to the evaluative and deontic realm in general, including moral norms but also values, theories, principles, views, etc.

Consequences are the most obvious example of how a lack of empirical information might inflict uncertainty about the moral value of an action (Crutchfield et al., 2023; Fleischhut et al., 2017; Kortenkamp & Moore, 2014; Meder et al., 2019; Shou et al., 2020; Shou & Song, 2017). Yet, researchers have long shown that facts about events leading up to an action can also be morally relevant, such as the intention of the agent (e.g., Greene et al., 2009; Monroe & Malle, 2019; Young & Saxe, 2009) or the broader context of the action (Andrejević et al., 2020; Schein, 2020), so that uncertainty about these facts can also result in moral uncertainty (Hansson, 2013; Tarsney, 2018).

We defined two requirements to consider someone empirically uncertain in the morally relevant sense: The person must be uncertain about empirical facts, and this uncertainty must reduce the certainty of their moral judgment. These two aspects need to be considered when trying to measure if someone is morally uncertain due to unknown empirical facts. One approach that focuses on the first aspect is asking participants for probability estimates relating to the act in question, such as how probable they think a specific outcome would be. This has been done, for example, by Kortenkamp and Moore (2014), Shou et al. (2020), or Shou and Song (2017). This approach can be assumed to be easily understood by study participants and it gives subjective probability estimates for various, clearly defined aspects of the described act. However, it neglects whether uncertainty regarding these outcomes negatively affects the person's overall moral certainty. Further, it cannot be guaranteed that the probability estimates researchers ask for exhaustively capture all empirical aspects that participants care about in a given situation. Lastly, the described approach requires adaptation for each scenario and thus lacks comparability across different contexts. However, such comparability is crucial, for example, when comparing across situations which uncertain facts lead to moral uncertainty.

A more general approach to measuring the degree of how much empirical information someone has to base their attitude on is described by the construct of subjective attituderelevant knowledge (Biek et al., 1996; Philipp-Muller et al., 2020; Wood et al., 1995). To assess this construct, participants are asked to indicate how much they know about the respective topic, for example, by asking "How much do you feel you know about mandatory vaccinations?" (Philipp-Muller et al., 2020). It is assumed that such knowledge contributes to the stability of one's respective attitude (Wood et al., 1995). Similar to empirical uncertainty, subjective attitude-relevant knowledge addresses the informational foundation of the respective attitude, thus both constructs are closely related. In contrast to empirical uncertainty in the morally relevant sense, however, subjective attitude-relevant knowledge does not focus on whether the specified level of knowledge affects the certainty of the moral judgment. It seems plausible that someone can have very little knowledge regarding a specific issue but still have high confidence in their moral judgment of it. Similarly, one can be highly informed about topics related to mandatory vaccinations, and still be morally uncertain due to some empirical aspects remaining unclear. This discrepancy would not be captured with this construct but is central to empirical uncertainty in the context of moral uncertainty.

In conclusion, we understand empirical uncertainty as moral uncertainty arising from ignorance of empirical facts that the respective individual deems relevant to their moral judgment. So far, research has focused on decision-making under empirical uncertainty in the context of uncertain consequences, but the construct can also refer to the circumstances of the specific act. When it comes to measuring empirical uncertainty, most existing approaches do not consider whether the described lack of empirical knowledge is actually relevant for the certitude of the moral judgment. We address this issue by proposing scale items that directly capture whether participants are morally uncertain due to the empirical facts.

Uncertain Values: Normative Uncertainty in Moral Judgments

When it comes to difficult moral cases, people might not only be uncertain due to unknown empirical facts. They might also be uncertain how to evaluate these facts from a moral perspective. Different conceptualizations of such normative uncertainty have been proposed in the philosophical literature. These conceptualizations largely differ in their metaethical assumptions about the nature of morality and moral practice (for an overview, see Theisen, 2023). Some philosophers understand the object of normative uncertainty to be putative moral facts (MacAskill et al., 2020; Tarsney, 2021), whereas others understand normative uncertainty as arising from the difficulty of applying abstract, individual moral values to concrete actions (Eriksson & Francén Olinder, 2016; Moller, 2011; Rosenthal, 2019). All of these approaches share the assumption that normative uncertainty can refer to one moral value by itself. Other approaches understand normative uncertainty as the result of conflict between multiple values, as they usually occur in moral dilemmas (Makins, 2021; Robinson & Steele, 2022). In this paper, we try not to take a stance on metaethical matters. We base our argument on a minimal definition of normative uncertainty as uncertainty about the normative evaluation of empirical facts, regardless of the metaethical nature of these evaluations.

Moral dilemmas are prototypical for cases that can elicit normative uncertainty (Robinson & Steele, 2022). Due to their relevance for the dual process theory, they have attracted considerable attention in the moral psychology literature; however, mostly without focusing on the resulting uncertainty (Ellemers et al., 2019; Greene et al., 2001; but see Białek & Neys, 2017; Mata, 2019; Skovgaard-Olsen & Klauer, 2024). There is also research on value conflict that goes beyond the sacrificial dilemmas used in dual process theory research, focusing on situations where aspects other than obligations and consequences have to be weighed against each other (Hanselmann & Tanner, 2008; Mandel & Vartanian, 2008; Shaddy et al., 2021). This research has shown that decisions involving conflict between moral values are perceived as more difficult than decisions involving conflict between non-moral values (Hanselmann & Tanner, 2008) and that different resolving strategies might be deployed depending on the type of trade-off that the situation demands (Shaddy et al., 2021).

While normative uncertainty is sometimes reduced to value conflict, philosophers have also discussed other constellations that might lead to normative uncertainty (Hicks, 2018; MacAskill, 2016; Moller, 2011; Weatherson, 2019). In those examples, normative uncertainty
arises from being confronted with ethical arguments that challenge one's previously held moral convictions, often by aiming at an extension of one's moral considerations, such as the inclusion of non-human animals or the introduction of higher moral demands than what is usually expected. The uncertainty in such cases is assumed to refer to the soundness of the respective argument. At first glance, this might seem like a phenomenon that concerns academic philosophers but not philosophical lay people. Yet, people are confronted with ethical arguments in everyday life, for example, by media consumption or exposure to political activism (e.g., Fernández, 2021; Tamborini et al., 2018). Even though the related process of attitude moralization has already been investigated empirically (Feinberg et al., 2019; Rhee et al., 2019), the normative uncertainty that might accompany the development of a moralized attitude so far has not received due attention in the psychological literature.

Empirical uncertainty can be resolved by learning the respective missing information (Alvarez & Brehm, 1995, 1997). Resolving normative uncertainty might not be as straightforward. Even though some philosophers assume that resolving normative uncertainty merely requires learning the respective moral fact and that it is thus, in principle, similarly resolvable as empirical uncertainty (e.g., MacAskill et al., 2020), a substantial share of philosophers have rejected this view (Bourget & Chalmers, 2023) and it is even less common among lay people (Theisen, 2023). Still, other ways have been described in which people might try to reduce their uncertainty about normative evaluations. People might adhere to the opinion of moral experts (Costa-Gomes & Schoenegger, 2023), they may consult with peers (Rowland, 2021), or they may rely on introspective processes, such as intuition, to discover what their values in the respective situation dictate (Theisen, 2023). If their normative uncertainty stems from conflicting values, they might also try to find empirical information that tips the scale in favor of one value over another (van Harreveld et al., 2009) – however, this assumes that the value conflict dissolves when more empirical facts are known, which might not always be the case.

There has been very limited empirical research explicitly addressing normative uncertainty and its measurement (Costa-Gomes & Schoenegger, 2023; Dietrich et al., 2019; Jabarian, 2020; Theisen, 2023). Yet, some studies have addressed very related concepts, such as overall moral certainty or value conflict (e.g., Alsaad, 2021; Alsaad et al., 2021; Hanselmann & Tanner, 2008). Usually, these studies do not differentiate between empirical and normative aspects of uncertainty. One exception to this is the work of Alvarez and Brehm (1995, 1997). They were not interested in moral judgment but in explaining response variability in political attitudes toward abortion and racial policies. They argued that variability underlying a person's response can be attributed, among other sources, to a lack of informedness or to value conflict, which partly maps the distinction of empirical and normative uncertainty. If the variability is due to a lack of informedness, they argued, it should correlate negatively with the amount of information the respective person has. If instead the variability correlates positively with value conflict, they interpret this to be indicative of ambivalence. Alvarez and Brehm's approach supports our argument that lacking attitudinal consistency might not only arise due to unknown empirical facts but also due to inconsistencies in the normative evaluation of the facts. Still, their approach cannot be considered a direct test of our hypothesis of empirical and normative uncertainty as two different psychological states: First, they were not interested in moral uncertainty but in response variability in the political context. Response variability may be one consequence of moral uncertainty; however, these constructs cannot be equated. Second, Alvarez and Brehm interpreted response variability resulting from value conflict as ambivalence, which poses the question of whether that equates to normative uncertainty. As we have shown above, both concepts can refer to value conflict; however, normative uncertainty is a broader term that can also refer to uncertainty about ethical arguments and individual values, which was not considered by Alvarez and Brehm and, as far as we understand it, would not be identifiable by their approach. Lastly, concerns regarding the reliability of Alvarez and Brehm's modeling technique have been voiced, since the model they propose was not able to adequately recover

simulated parameters (Keele & Park, 2006). Despite these limitations, the work by Alvarez and Brehm shows that dissecting different aspects of attitudinal uncertainty might be worthwhile. Instead of relying on the modeling of response variability, we chose a primarily factor-analytical approach to differentiate between different forms of moral uncertainty.

The Current Research

In line with recent philosophical work on moral uncertainty (Lockhart, 2000; MacAskill et al., 2020; Nissan-Rozen, 2015), we argue that empirical and normative uncertainty are two psychologically distinguishable states that require separate consideration. We address this issue by introducing the Moral Uncertainty Scale (MUS), which reliably and validly measures empirical and normative uncertainty across a wide range of moral situations. In Study 1, the scale items were developed and the underlying factor structure was identified. This was done on a dataset where each participant self-reported a situation that they were morally uncertain about. In Study 2, we combined experimental and factor-analytical approaches to confirm the identified factor structure on a set of predefined scenarios. All studies reported in this paper were conducted in line with the Declaration of Helsinki. Informed consent was obtained before each data collection.

Study 1: Scale Development

The aim of Study 1 was to develop and analyze a pool of items that would maximally differentiate between empirical and normative uncertainty. We started with the formulation of items based on the philosophical literature. These were then subject to expert ratings by philosophers and cognitive interviews with non-philosophers to ensure expert and face validity. In a next step, we developed a method that prompted participants to describe a scenario of moral uncertainty. This method allowed us to investigate the item pool without the need to define a specific moral scenario that the items refer to, leading to a greater variety of stimulus material. This ensured that the final scale would not depend on the characteristics of a specific scenario it was developed on and ensured external validity (Bauman et al., 2014). We used an exploratory factor-analysis (EFA) to identify and extract the factors underlying our item pool. Graded response modeling was used to assess the quality of the items, to reduce the length of the factor scales, and to select the most fitting response format.

Method

Item Pool

Item Wording. Following our definition of empirical uncertainty, the respective items described being morally uncertain due to a lack of unknown circumstances or unclear consequences (e.g., "The moral evaluation of the action is difficult because some facts are unclear", "It is difficult to judge the action because I do not know what doing or not doing the action would lead to"). Items for normative uncertainty were derived from the philosophical literature on decision-making under normative uncertainty (e.g., MacAskill et al., 2020; Moller, 2011; Nissan-Rozen, 2015) and on the metaethics of normative uncertainty (e.g., Eriksson & Francén Olinder, 2016; Makins, 2021; Robinson & Steele, 2022), including aspects such as uncertainty regarding the right moral view and conflict between opposing moral views (e.g., "I am unsure how to apply my moral concepts to this situation", "I have different moral views that lead to different evaluations of this action"). The items were developed in English.

Expert Interviews. The validity of these items was rated by a sample of professional philosophers (*N* = 9) who had published peer-reviewed work on normative uncertainty before. Items with a high validity rating were kept in the item pool. All negatively framed items were given a substantially lower rating than the positively framed items. It was not clear whether this was due to a misunderstanding or their poorer validity. These items were kept, regardless of their poor rating, for further analyses. One expert voiced concern that some items would refer to the absence of moral expertise instead of moral uncertainty. However, as other experts gave the respective items rather high validity ratings, we decided to keep them in the item pool. Others noted that the term *fact* was not specified as *empirical fact* (as opposed to *moral fact*) in the items. This was a point of focus in the following cognitive interviews.

Cognitive Interviews. We first conducted *N* = 7 cognitive interviews with the initial English items. Participants were recruited on the campus and in the private environment of one of the authors. Participants, who were all non-philosophers, read a dummy scenario about abortion and prenatal diagnosis, after which they were asked to voice their thoughts while responding to the item pool. Ambiguous or misleading items were edited after each interview, so that each participant was presented with a slightly updated version of the item pool. The latter two of these interviews did not lead to further fundamental adaptations. The issue of empirical facts only being described as *facts* did not come up in any of the interviews; the term *facts* was consistently understood as intended. Following this initial phase of cognitive interviews, we decided that the subsequent steps of the scale development would have to be conducted in a German convenience sample due to financial limitations. Therefore, we translated the most recent version of the items to German. These translations were then subject to three further cognitive interviews which led to only minimal further adaptations of the item pool.

Pre-test of Scenario Generation

To ensure that we would not identify a factor structure that is unique to a specific moral scenario but that would be robust across a large variety of stimulus material, we asked participants to describe a situation that they have experienced in their own life and where they were now wondering what would have been the morally right thing to do. Consequently, each participant rated the items for a different scenario. This procedure was inspired by Gerpott et al. (2018) who used a similar approach to develop a scale of subjective outcome interdependence in a specific situation. We conducted a pre-test on N = 50 undergraduate students, recruited via a mailing list from the Department of Psychology at Heidelberg University, Germany, to test whether participants were able to generate adequate scenarios. Participants received partial course credit for their participation. They were instructed that they might describe situations from personal or professional contexts, situations of minor or major moral relevance, and

situations that they had experienced for themselves or that they have merely been told about. It was stressed that their description must feature a specific action of whose moral evaluation they were now uncertain. As an illustration, some example situations were given, such as "Deciding whether it would have been right for a friend to visit their family despite the risk of infecting them with a virus" or "Deciding whether it would be right to assist a struggling classmate in cheating on an exam to prevent them from failing". It was pointed out that these were just examples and that their own situations might vary drastically.

One participant's response was excluded from the analyses as they skipped the part where they were asked to describe a situation. From the remaining 49 participants, 47 described situations where they were uncertain about the moral value of a specific action. We took this as indicative that our task instructions successfully prompted participants to describe situations of moral uncertainty.

Factor Extraction and Item Selection

Participants. For an EFA, we collected data from N = 332 participants. Participants were recruited via mailing lists of several German universities. Six participants indicated that they had not responded conscientiously. Further 13 participants had a relative speed index > 2, indicating that they may have rushed the assignment (Leiner, 2019). From the remaining sample, 26 participants described cases that either did not show any moral relevance or where no uncertainty was involved. A further five participants had a mean item score < 1.5, which indicated that they almost exclusively chose the lowest rating. Finally, 17 participants were excluded because they had more than four missing values. The final sample comprised of N =265 participants (197 female, 64 male, 2 non-binary, 2 not indicating) aged between 18 and 70 (M = 24.67, SD = 6.9). As Mundfrom et al. (2005) suggest, our sample size should be sufficient to identify up to six factors with low levels of communality assuming that we have at least 6 items per factor. The sample size was thus suited for our aim. As reimbursement, participants received partial course credit or could join a lottery for two vouchers of 25€ value each. **Procedure.** The study was conducted online using *SoSci Survey* software (Leiner, 2022). First, similar to the pre-test, we asked participants to describe a situation involving a morally ambiguous action. They were then asked to report their uncertainty regarding the moral evaluation of the described action by rating the items from our pool on a seven-point Likertscale.² The items were presented in a randomized order. Participants were told to focus on moral aspects and that their moral evaluation might differ from what they might have actually done or would actually do in the described situation.

Having responded to our item pool, we additionally asked participants for a first-order, binary judgment of whether they think it would be morally better to do or not to do the action they had described. Then, they were asked how certain they were in their first-order judgment on a slider scale ranging from 0 (*not certain at all*) to 100 (*completely certain*). This measurement was included to capture overall moral certainty. One session took approximately 20 minutes.

Analysis. All statistical analyses were conducted in *R* (version 4.1.2; R Core Team, 2020). After participants with more than four missing values had been excluded, the remaining missing values were imputed using the *mice* algorithm (version 3.13.0; van Buuren & Groothuis-Oudshoorn, 2011). The dataset was then submitted to an EFA using the *psych*-package (version 2.1.9; Revelle, 2021). Principal axis extraction method was selected (Watkins, 2018). To account for the ordinal nature of our data, we used polychoric correlations. The number of factors was determined based on parallel analysis (Horn, 1965). Because we assumed correlated factors, a *promax* rotation was employed (Hendrickson & White, 1964).

Items were removed from the dataset if their highest factor loading was < .4, if they showed relevant cross-loadings (i.e., factor loadings >= .4 for more than one factor), or if they belonged to a factor with fewer than three items or with an internal consistency < .7 (Watkins, 2018). EFA was repeated after each round of item removal.

² The full list of items can be found in Supplementary Table S1.

Items of the thus identified factors were further analyzed in graded response models (Samejima, 1969) using the *mirt*-package (version 1.37.1; Chalmers, 2012). Graded response models are established models in the item response theory context. They are suited for polytomous items, such as the Likert-type items used in this study (Samejima, 1969). These models were used to select the most informative items for each factor and thereby reduce the overall length of the scale.

Lastly, we conducted a confirmatory factor analysis (CFA) on the final item selection to estimate latent correlations between the identified factors and the measurement of overall moral certainty. This was done using the *lavaan*-package (version 0.6-14; Rosseel, 2012).

Results

Data Screening

After the exclusion of participants with more than 4 missing values, 187 (1.68%) missing values remained in the dataset, distributed across 99 participants. Little's MCAR test (Little, 1988) was not significant, χ^2 (2907) = 1353, p > .99, indicating that missing values occurred at chance. Therefore, missing values could be imputed. Skewness and kurtosis were investigated for each item in the pool. All values were within -2 and 2, indicating that none of the items deviated excessively from normal distribution. However, the boxplot of one item (NU19: "I know which aspects of the situation are morally relevant") revealed a skewed distribution so that any value above 4 was marked as an outlier. Based on this observation, we decided to exclude this item from further analyses.

Which Factors Could Be Extracted From the Dataset?

The remaining 40 items were submitted to an EFA to determine their factor structure. Bartlett's test of sphericity was significant, $\chi^2(780) = 5219.181$, p < .001, indicating that the correlation matrix significantly differed from an identity matrix (Bartlett, 1950), and the Kaiser-Meyer-Olkin criterion was > .7, *KMO* = .911, suggesting that the data was suited for factor analysis (Kaiser, 1974; Watkins, 2018).

Parallel analysis indicated a five-factor solution. In the subsequent factor analysis, four items did not load on any of the factors and were excluded from further analyses. Parallel analysis on the now reduced dataset still indicated a five-factor solution and the factor analysis was repeated, this time without problematic loadings. Except for one item (NU06: "I am unsure how to evaluate the consequences of doing or not doing the action from a moral point of view"), each factor only consisted of items from either the empirical or the normative uncertainty section, so that a clear distinction between both dimensions emerged. Based on the semantic content of the items, the two empirical uncertainty factors could be described as "lack of information" (e.g., "One would need more information to be able to judge the situation clearly"; α = .913) and "unclear consequences" (e.g., "The moral evaluation of the action is so difficult because some consequences are unclear"; $\alpha = .871$). The three factors comprising of items aiming at normative uncertainty were "normative uncertainty" (e.g., "Conflicting arguments come to my mind when judging this action"; α = .873), "moral confidence" (e.g., "Even if I had doubts about some of the facts, my moral assessment of the situation is clear"; α = .688), and "moral relevance" (e.g., "I am unsure whether doing or not doing the act is a question of morality"; $\alpha = .696$). Due to their low internal consistency, we dropped the items for "moral confidence" and "moral relevance", following our a priori defined criteria (Watkins, 2018). For the remaining items, parallel analysis suggested a three-factor solution. A subsequent factor analysis again revealed the factors "lack of information" (accounting for 19% of total variance and 36% of common variance after rotation), "unclear consequences" (16% of total variance and 32% of common variance), and "normative uncertainty" (17% of total variance and 32% of common variance), comprising of the same items as before (Table 2). "Lack of information" correlated positively with "unclear consequences", r = .640, and with "normative uncertainty", r = .471. "Unclear consequences" and "normative uncertainty" showed a correlation of r = .435.

We used graded response modeling (Samejima, 1969) to reduce the number of items per factor and to determine the appropriate number of response options. A detailed report of

these analyses can be found in the Appendix. The "lack of information" subscale could be reduced to five items, the "unclear consequences" subscale could be reduced to four items, and the "normative uncertainty" subscale could be reduced to nine items. Analyses of trace lines suggested a six-point Likert-scale as the response format.

A CFA on the same dataset, which included the described factors after their reduction and the single-item measure of overall uncertainty, revealed significant correlations between overall moral certainty and "lack of information", r = -.307, "unclear consequences", r = -.287, and "normative uncertainty", r = -.559, *CFI*_{robust} = 0.951, *TLI*_{robust} = 0.943, *RMSEA*_{robust} = .060 90% CI [.047; .072], *SRMR* = 0.045.³ When introducing a second-order factor for "lack of information" and "unclear consequences" for describing empirical uncertainty, this factor correlated with overall moral certainty with r = -.346.

Discussion

In Study 1, we developed an item pool for measuring empirical and normative uncertainty. The items were validated by ratings from expert philosophers and cognitive interviews with non-philosophers. An initial EFA suggested five factors of which only three had sufficient reliability. However, even though they were excluded from further analyses, the omitted factors "moral confidence" and "moral relevance" should be briefly discussed, since the reliability of each was close to the a priori defined threshold. "Moral confidence" comprised of items capturing certainty in one's judgment. In hindsight, these items were not suited to distinguish between empirical and normative sources of uncertainty: If people scored low on them, this could be due to empirical uncertainty, normative uncertainty, or both. They were therefore not suited for our goal. "Moral relevance" comprised of two items referring to whether the questionable act has any moral relevance at all and two items referring to the moral relevance of individual aspects of the act. Whilst the latter matched one of the theoretical

³ The identified factor structure was further confirmed on a separate dataset in Study 2.

facets of normative uncertainty, namely uncertainty regarding an individual value, the number of items was not sufficient to consider this as an individual factor. It should also be noted that the initial item pool comprised of further items with similar content (e.g., "With some aspects of the situation I am unsure whether I should consider them in my moral judgment", "Some aspects of the situation make me wonder if they have any bearing on how good or bad the action is morally"); however, these could not be assigned clearly to any factor, which further obscured the interpretation of the "moral relevance" factor. The difficulty in interpreting the content of the "moral relevance" factor further supported our decision not to consider it in the further scale development.

The remaining items formed three reliable subscales, two of which addressed empirical uncertainty ("lack of information", "unclear consequences") and one that addressed normative uncertainty ("normative uncertainty"). All three factors correlated negatively with overall moral certainty; the highest correlation could be shown for the "normative uncertainty" factor.

In the subsequent analyses, we treated "lack of information" and "unclear consequences" as subfactors of a second-order factor "empirical uncertainty" due to theoretical considerations, due to them only consisting of items that were written with empirical uncertainty in mind, and due to their high intercorrelation. Note that a test of a model with this second-order factor against a model without it was not possible because those models are statistically equivalent. Yet, the introduction of the second-order factor facilitated interpretation and contrasting of empirical and normative aspects.

Findings from our first study supported the assumption that uncertainty due to ignorance of empirical facts, be it about context or consequences, and normative uncertainty are distinct psychological states. These findings were further validated in Study 2.

Study 2: Scale Validation

In Study 2, we aimed to validate the factors that we had identified in Study 1 on a separate dataset. In addition to a factor-analytical confirmation, we also included experimental

manipulations of the stimulus material. We asked participants to rate eight scenarios that varied in setting, certainty of information (targeting empirical uncertainty), and value conflict (targeting normative uncertainty). We also used this stimulus material to show that the identified factor structure is not specific for self-generated scenarios but is robust and invariant across predefined stimuli as they are commonly used in psychological research. Participants rated the scenarios on the MUS items, on four items capturing overall moral certainty (O'Connor, 1995), and on the clarity/correctness items by Petrocelli et al. (2007). The latter scale describes two dimensions of attitude certainty: clarity regarding what one's true attitude toward the object is, and the sense that one's attitude is the correct attitude to have toward that object. We used this scale to ensure that the distinction between empirical and normative uncertainty is different from the established distinction between clarity and correctness.

The hypotheses for Study 2 were preregistered on the *Open Science Framework* in advance.⁴ In a first step, we tested the factor structure of the three subscales and their intercorrelations in a CFA. We hypothesized that a respective measurement model would show acceptable fit to the data (*H1*). This included the assumption of a higher-order factor "empirical uncertainty" comprising of "lack of information" and "unclear consequences", and a positive correlation between this "empirical uncertainty" factor and the "normative uncertainty" factor (*H2*). We assumed that this measurement model would be invariant across different scenario configurations (*H3*). Second, we were interested in the subscales' correlations with overall moral certainty, which we considered an indicator of convergent validity (*H4*). Third, for discriminant validity, we addressed overlap and differences between the MUS subscales and the two already established dimensions of attitude certainty, which are clarity and correctness (Petrocelli et al., 2007; Philipp-Muller et al., 2020). We hypothesized that empirical and normative uncertainty are distinct from clarity and correctness (*H5*), and that empirical and normative uncertainty each

⁴ https://osf.io/hb5mz/

contribute incrementally to a model that predicts overall moral certainty based on clarity and correctness (*H6*). Last, we experimentally manipulated empirical and normative uncertainty to ensure that the MUS subscales differentiate between both constructs. We hypothesized that empirical uncertainty would be higher in scenarios where some information is explicitly unknown in contrast to scenarios where information is described as definite (*H7*). We further hypothesized that normative uncertainty would be higher in scenarios where information scenarios where there is a value conflict in contrast to scenarios where the conflict is merely between different means to reach the same end (*H8*).

Method

Participants

We aimed for a sample size of N = 400. Participants were recruited in two rounds. In the first round, we recruited 146 participants (Table 3) from mailing lists of German universities. They received course credit for their participation. A further 256 participants were recruited via *Prolific*, who received £6 for their participation. Participants were excluded from the analyses if they indicated that they had not responded conscientiously, if they had more than one failed attention check, or if they had more than 10% missing values. In total, 482 participants started the study, 432 finished it, and 402 fulfilled the inclusion criteria, which then formed our final sample.

Scenarios

The scenarios were developed based on literature on tragic trade-offs (Hanselmann & Tanner, 2008; Mandel & Vartanian, 2008; Shaddy et al., 2021) and on studies on the role of empirical uncertainty in moral judgments (Kortenkamp & Moore, 2014; Shou et al., 2020). Each scenario described a situation with a choice between two actions. The scenarios ended with the question of whether it would be morally permissible to choose action A over action B. The scenarios had two different settings. In the first setting, there were four scenarios that featured a noil ship in distress. In the second setting, there were four scenarios that featured a company

that has to react to public criticism. In each setting, there was a baseline scenario where both empirical and normative uncertainty were expected to be low. In this baseline scenario, the consequences of both options were clearly stated and both options tried to achieve the same aim, however one of them was described as known to be better than the other (e.g., choosing a climate protection policy with known efficacy of 45% vs. a climate protection policy with known efficacy of 15%). Then, there was a scenario with high empirical uncertainty and low normative uncertainty where the effectiveness of one of the options was unclear so that there was uncertainty about which of the two options was better suited to achieve the same aim (manipulation of empirical uncertainty; e.g., choosing a climate protection policy with known efficacy of 15% vs. a climate protection policy with questionable efficacy of 40%). In a further scenario variant, there was low empirical uncertainty and high normative uncertainty. In this variant, the consequences of both options were stated clearly. However, both options would lead to conflicting goals, creating a tragic trade-off/dilemma (manipulation of normative uncertainty; e.g., choosing a climate protection policy with known efficacy vs. an occupational safety policy with known efficacy). Lastly, with high empirical and high normative uncertainty, one scenario in each setting featured two options with conflicting aims and with uncertainty regarding the efficacy of one of these options to reach the respective aim (interaction of empirical and normative uncertainty; e.g., choosing a climate protection policy with questionable efficacy vs. an occupational safety policy with known efficacy). The full scenario descriptions can be found in the preregistration.

Procedure

The study was conducted online using *SoSci Survey* software (Leiner, 2022). For each scenario, the following measurements were taken: moral judgment on a scale from 1 (*morally absolutely wrong*) to 6 (*morally absolutely right*); four items derived from O'Connor (1995) measuring overall moral certainty regarding this judgment on a seven-point Likert-scale; the clarity and correctness items by Petrocelli et al. (2007) using a seven-point Likert-scale; and lastly

the MUS items in a randomized order on a six-point Likert-scale ranging from 1 (*not at all true*) to 6 (*totally true*). The overall moral certainty scale and the clarity/correctness scale were translated to German by the authors. The study also featured three attention check items, two of which appeared among the MUS items and one among the clarity/correctness items. Median duration time was 25.7 minutes.

Analysis

Hypotheses regarding factor structure and intercorrelations were tested via structural equation modeling using the *lavaan*-package (version 0.6-14; Rosseel, 2012) in *R* (version 4.1.2; R Core Team, 2020). We started by assessing the fit of the respective measurement models before modeling associations between latent variables. Because of the repeated measurements design (each participant rated eight scenarios), we fit models with cluster-robust standard errors. Since all variables that were included in our models were on the trial-level, these cluster-robust standard errors were suited to account for the grouping of observations within participants in our data (Hazlett & Wainstein, 2022; McNeish et al., 2017; see also Reise et al., 2005). As robustness checks, we had also intended to fit full multilevel models. However, these only converged after several modifications to the models, which undermined comparability. We therefore decided to directly interpret the models based on cluster-robust standard errors instead.

We tested measurement invariance across scenarios from different scenario settings (oil vs. company), across different levels of the empirical uncertainty manipulation (low vs. high), and across different levels of the normative uncertainty manipulation (low vs. high). Due to the inclusion of the second-order factor for "empirical uncertainty", we followed the procedure described by Rudnev et al. (2018) for assessing measurement invariance with second-order factors. We tested for configural (same model syntax in both subgroups), metric (same loadings in both subgroups, both on first order and second order level), scalar (same loadings and

intercepts in both subgroups, both on first order and second order level), and strict measurement invariance (same loadings, intercepts, and residuals in both subgroups).

For testing hypotheses *H6-H8*, we fitted linear mixed models with random intercepts for each person using the *Ime4*-package (version 1.1-31; Douglas Bates et al., 2015) in *R*. Random slopes were included if they further improved model fit. For the linear mixed models, factor scores were calculated by weighting item scores with their corresponding loadings from the standardized solution of a structural equation model that included the MUS factors, overall certainty, clarity, and correctness, as well as their intercorrelations (Figure 1).

Results

Data Screening

There were no missing values for any of the scale items. There were numerous cases where a participant gave the same response to all items of a scale within a given scenario. There were also participants who gave the same response to a specific item across all scenarios, which was a problem for full multilevel modeling. As preregistered, we report findings from the full dataset.⁵

Descriptive statistics for unweighted composite scores of the variables are shown in Table 4. Skewness and kurtosis values for each item were between -2 and 2, indicating that there was no substantial deviation from normal distribution. *ICC* for each item was > .05, indicating substantial correlation within participants.

Measurement Models

Hypotheses 1 and 2: Could the Assumed Factor Structure of the MUS Be Confirmed? A CFA with the factors "lack of information", "unclear consequences" and "normative uncertainty", and a second-order factor "empirical uncertainty" comprising of "lack of

⁵ The overall pattern of findings was robust across various data preparation and analysis decisions, including the removal of cases consisting of the same response for all items.

information" and "unclear consequences" showed good model fit, as indicated by *CFI*_{robust} >= 0.90 and *RMSEA*_{robust} <= .08, *CFI*_{robust} = 0.975, *TLI*_{robust} = 0.970, *RMSEA*_{robust} = .069 90% CI [.065; .072], *SRMR* = 0.026. Thereby, the factor structure that was derived from Study 1 could be confirmed (*H1*). Setting the correlation between empirical uncertainty and normative uncertainty to zero reduced model fit, *CFI*_{robust} = 0.933, *TLI*_{robust} = 0.921, *RMSEA*_{robust} = .111 90% CI [.107; .115], *SRMR* = 0.378, $\chi^2_{diff}(1) = 427.67$, p < .001, $\triangle AIC = -2517$. The hypothesis that empirical and normative uncertainty are correlated could therefore also be confirmed (*H2*). Analysis of the full model including all observed variables suggested a correlation of r = .781 (Figure 1).

Hypothesis 3: Is the MUS Invariant Across Different Scenario Manipulations? We

tested for measurement invariance across different scenario settings, different levels of the empirical uncertainty manipulation, and different levels of the normative uncertainty manipulation (Table 5). For all comparisons, we found strict measurement invariance, both for first- and second-order factors.

Fitting models for different levels of the empirical uncertainty manipulation led to negative estimates for the variance of the "lack of consequences" factor in scenarios where the empirical uncertainty manipulation was zero. However, these estimates were not significantly different from zero, suggesting that model misspecification could be ruled out (Kolenikov & Bollen, 2012). As a robustness check, we also tested invariance without the second-order factor, which further confirmed the strict invariance assumption. *H3* could be confirmed.

Associations Between Variables

Hypothesis 4: Do Empirical and Normative Uncertainty Correlate With Overall Moral Certainty? Analyses of the measurement model for overall moral certainty suggested a correlation between the two positively framed items of that factor. We included this correlation in the further analyses. After this adaptation, model fit of the measurement model was good, $CFI_{robust} = 1.00, TLI_{robust} = 0.998, RMSEA_{robust} = .034 90\%$ CI [.0; .086], SRMR = 0.002. To test our hypothesis that empirical and normative uncertainty correlate with overall moral certainty, we fitted a model containing the MUS factors and the overall moral certainty factor, as well as their correlations, *CFI*_{robust} = 0.976, *TLI*_{robust} = 0.973, *RMSEA*_{robust} = .058 90% CI [.055; .061], *SRMR* = 0.025. Fixing the correlation between empirical uncertainty and overall moral certainty to zero reduced model fit in comparison to the unrestricted model, *CFI*_{robust} = 0.952, *TLI*_{robust} = 0.945, *RMSEA*_{robust} = .083 90% CI [.080; .086], *SRMR* = 0.286, $\chi^2_{diff}(1) = 702.16$, *p* < .001, \triangle AIC = -1876. Similarly, fixing the correlation between normative uncertainty and overall moral certainty to zero reduced model fit in comparison to an unrestricted model, *CFI*_{robust} = 0.936, *TLI*_{robust} = 0.927, *RMSEA*_{robust} = .096 90% CI [.093; .099], *SRMR* = 0.287, $\chi^2_{diff}(1) = 677.66$, *p* < .001, $\triangle AIC = -3073$. Therefore, the hypothesis that empirical and normative uncertainty correlate with overall moral certainty could be confirmed (*H4*). Analysis of the full model including all observed variables suggested a correlation between normative uncertainty and overall moral certainty of *r* = -.708, and a correlation between normative uncertainty and overall moral certainty of *r* = -.829 (Figure 1).

Hypothesis 5: Are Empirical and Normative Uncertainty Distinct From Clarity and

Correctness? The measurement model for clarity and correctness had a close to acceptable fit, however, *RMSEA* was above the threshold of .08, *CFI*_{robust} = 0.977, *TLI*_{robust} = 0.962, *RMSEA*_{robust} = .117 90% CI [.105; .129], *SRMR* = 0.027. This could not be resolved by theoretically meaningful adaptations. Since internal consistency and factor loadings were high, we decided to proceed with the analyses without further adaptation.

To test discriminant validity, we followed the protocol described by Rönkkö and Cho (2022). We fitted a model containing the MUS factors and the factors clarity and correctness. The assessment of discriminant validity was based on the upper limit of the 95% confidence interval of the absolute correlations between the respective factors. If $r_{upperlimit} < 0.8$, discriminant validity was assumed (i.e., the corresponding factors were assumed to be empirically distinct). The upper limit for the absolute correlation between empirical uncertainty

and clarity was $r_{upperlimit} = 0.621$; for empirical uncertainty and correctness, it was $r_{upperlimit} = 0.701$. For the correlation between normative uncertainty and clarity, the upper limit was $r_{upperlimit} = 0.743$, and for normative uncertainty and correctness, it was $r_{upperlimit} = 0.792$. Since all these values were below the cutoff of $r_{upperlimit} = 0.8$, discriminant validity could be stated. This finding was further confirmed by testing the model against models where the respective correlations were restricted to $r \ge .8$ each at a time. For each investigated correlation, the restricted model had a significantly worse fit than the unrestricted model. Hence, although MUS and the clarity/correctness scales were highly correlated, they were not redundant (i.e., $r_{supperlimit} < 0.8$) and *H5* could be confirmed.

Hypothesis 6: Do the MUS Factors Explain Additional Variance in Overall Moral Certainty? To test whether empirical and normative uncertainty would incrementally explain variance in overall moral certainty in addition to the variance explained by clarity and correctness, we started by fitting a baseline linear mixed model regressing overall moral certainty to clarity and correctness, and their interaction. Confirming *H6a* (i.e., incremental validity of empirical uncertainty), adding the empirical uncertainty score improved model fit, $\chi^2_{diff}(1) = 197.75$, p < .001, $\triangle AIC = 195.7$. The final model resulting from a further inclusion of random slopes is described in Table 6.

In a second step, we added a fixed effect for the normative uncertainty score to the baseline model. Confirming H6b (i.e., incremental validity of normative uncertainty), this also improved model fit, $\chi^2_{diff}(1) = 497.04$, p < .001, $\triangle AIC = 495$. The final model resulting from a further inclusion of random slopes is described in Table 7.

To get a full picture, we also fitted a model including fixed effects for clarity, correctness, and their interaction, empirical uncertainty, and normative uncertainty. Again, confirming that empirical and normative uncertainty incrementally explain variance in overall moral certainty, this model had better fit than the models containing only empirical uncertainty, $\chi^2_{diff}(1) = 339.31$, p < .001, $\triangle AIC = 337.3$, or only normative uncertainty in addition to clarity and correctness, $\chi^2_{diff}(1) = 40.02, p < .001, \triangle AIC = 38$. Adding an interaction between empirical and normative uncertainty did not further improve model fit, $\chi^2_{diff}(1) = 0.6734, p = .412, \triangle AIC = -1.3$ (for the final model after the inclusion of further random slopes see Table 8).

Exploratory: Can High Correlations Between MUS Factors and Clarity/Correctness Be Explained by Overall Moral Certainty? Concerning discriminant validity, we observed that the MUS and the clarity/correctness scales were highly correlated but not redundant (i.e., *rsupperlimit* < 0.8). Concerning the incremental validity, we observed that models each explained approximately 80% of variance in overall moral certainty; however, the individual contribution of each scale was low, again reflecting their high intercorrelation. Consequently, we suspected that the high correlations between the MUS and the clarity/correctness scales were driven by the fact that they were all correlated with overall moral certainty. Hence, to receive an uninflated measure of the intercorrelations of these scales, we fitted an additional structural equation model to estimate the correlations between clarity/correctness and the MUS factors while controlling for overall moral certainty (Table 9). When controlling for certainty, MUS factors and clarity/correctness were only weakly correlated, underlining the discriminant validity of empirical and normative uncertainty versus clarity and correctness.

Hypotheses 7 and 8: Are the Empirical and Normative Uncertainty Factors Sensitive to

Their Respective Manipulation? We examined whether introducing uncertain information (*H7*) or conflicting values to a scenario (*H8*) affected the empirical and normative uncertainty factors, respectively.⁶ For each MUS factor, we fitted several linear mixed models including fixed effects for both manipulations and their interaction. Both models benefitted from the addition of a fixed effect for scenario setting. The final models, including random slopes, are described in Table 10. In both models, the manipulation corresponding to each uncertainty factor had a

⁶ Similar patterns emerged when predicting lack of information or unclear consequences separately, as reported in the Supplementary Tables S2–S5.

greater effect on its corresponding factor score than the opposite manipulation. Hence, *H7* and *H8* were confirmed.

Interestingly, in both models, there were strong interaction effects. Even though empirical and normative uncertainty scores increased when the respective uncertainty was experimentally induced, these main effects were mostly driven by comparing cases where the other manipulation was not present (Figure 2). This interaction pattern could be observed both when predicting the empirical uncertainty score as well as the normative uncertainty score. We suspected that this interaction pattern might be explained by something that both uncertainty scores have in common – namely, their shared variance with overall moral certainty. More precisely, we suspected that the observed interaction effects were expressions of the interaction effect that both manipulations had on overall moral certainty. Following this line of reasoning, we fitted a further model predicting overall moral certainty from the manipulations. The respective model is described in Table 11. These analyses revealed a similar pattern of main and interaction effects as the models predicting empirical and normative uncertainty scores, indicating that the observed interaction patterns might be explained by changes in overall moral certainty.

Having established that the observed interaction effects might be attributed to shared variance with overall moral certainty, we reran the models predicting empirical and normative uncertainty scores with added fixed effects for overall moral certainty (Table 12). Controlling for overall moral certainty further improved model fit for predicting empirical uncertainty, $\chi^2_{diff}(1) = 1442.5$, p < .001, $\triangle AIC = 1440.5$, and for predicting normative uncertainty, $\chi^2_{diff}(1) = 2377.4$, p < .001, $\triangle AIC = 2375.4$. The empirical uncertainty score was now only affected by the empirical uncertainty manipulation and a negligible interaction effect between both manipulations. The normative uncertainty score was mainly affected by the normative uncertainty manipulation and to a lesser degree by the empirical uncertainty manipulation, while the relevance of the interaction term was greatly reduced. Controlling for overall moral

certainty greatly increased the specificity of the manipulations' effects on the respective MUS factors (Figure 3).

Discussion

In Study 2, we validated the factor structure of the MUS on a second dataset. We considered various scenario settings and different levels of empirical and normative uncertainty. Our preregistered hypotheses could all be confirmed. We found that the respective items loaded highly onto their corresponding factors. The model showed strict measurement invariance between "company" and "oil" scenarios, between different levels of the empirical uncertainty manipulation, and between different levels of the normative uncertainty manipulation.

The MUS factors correlated highly with clarity and correctness. Yet, we could state discriminant validity. Additionally, we could show that adding empirical or normative uncertainty to a model predicting overall moral certainty from clarity and correctness contributed incrementally to the explained variance.

Even though our criteria for stating discriminant validity were met, the data suggested a strong overlap between the MUS factors and clarity/correctness. Exploratorily, we found that this overlap was due to shared variance explained by overall moral certainty. Consequently, after controlling for overall moral certainty, correlations between MUS factors and clarity/correctness were drastically reduced. Additionally, both empirical and normative uncertainty showed stronger associations with correctness than with clarity. This suggests that neither empirical nor normative uncertainty can be clearly mapped to one of the clarity/correctness dimensions but that they describe two different distinctions of uncertainty.

The correlations between the MUS and the clarity/correctness dimensions might further inform us about the nature of empirical and normative uncertainty themselves. From a theoretical point of view, one might have expected that empirical uncertainty would be correlated with the correctness dimension, since empirical uncertainty refers to statements about facts that can be objectively correct or incorrect. Normative uncertainty, on the other side, was defined as referring to unclear or conflicting values, which might have been expected to resemble the clarity dimension more closely. Yet, the association between normative uncertainty and correctness was stronger than that between normative uncertainty and clarity. This also counters the view that normative uncertainty is reducible to ambivalence, since it suggests that normative uncertainty is associated with a lack of certainty about which moral view is correct, which goes beyond the concept of ambivalence as it is discussed in metaethics (Makins, 2021; Theisen, 2023).

To further validate our interpretation of the MUS factors, we tested how these factors would be affected by various modifications to the presented scenarios. We found significant main effects for the introduction of uncertain information on the empirical uncertainty factor and for the introduction of value conflict on the normative uncertainty factor. To a lesser degree, we also found that the empirical uncertainty manipulation increased the normative uncertainty score and vice versa. After controlling for overall moral certainty, the effect of the normative uncertainty manipulation on the empirical uncertainty score was greatly reduced. These findings suggest that the factors were sensitive towards the respective manipulations. We also found strong interaction effects between the empirical and normative uncertainty manipulations. Since similar interaction effects were also present when predicting overall moral certainty from the manipulations, it seems plausible to assume that these were not inherent to the MUS but rather to the scenarios themselves. Overall, our findings suggested two forms of interplay between empirical and normative uncertainty. First, they suggested that a lack of empirical information might also lead to normative uncertainty. This is in line with Nissan-Rozen (2015) who argues that unresolved empirical uncertainty poses the question of what is the morally right way to act under empirical uncertainty, which then is a matter of normative uncertainty. Second, the positive effect of the interaction of empirical and normative uncertainty on overall moral certainty suggests that adding uncertain information to a case of value conflict or adding value conflict to a case of uncertain information does not necessarily

increase the moral uncertainty any further. This is in line with previous findings that uncertain information can reduce the impact of value conflict and vice versa (Reich & Wheeler, 2016; van Harreveld et al., 2009).

Study 2 also revealed a high correlation between normative uncertainty and overall moral certainty, which might indicate marginal problems with discriminant validity between these constructs (r = -.829). The high correlation was surprising given a correlation of r = -.559 in Study 1. There are two possible explanations for this. The first concerns the difference in the measurement of overall moral certainty in both studies (single-item measure in Study 1 and four-item measure in Study 2). It may be possible that the certainty items that were used in Study 2 more closely matched normative uncertainty than the single-item measure in Study 1; however, this would not explain the also rather high correlation between overall moral certainty and empirical uncertainty. The second possible explanation concerns the difference in the rated scenarios. It may be that there was more variance in self-generated scenarios that could not be captured by these scales, which would have reduced the identified correlation. Nevertheless, the high correlation between normative uncertainty and overall moral certainty suggests that uncertainty in moral issues might largely refer to non-empirical aspects.

Due to statistical equivalence, it was not possible to determine whether introducing a second order factor "empirical uncertainty" comprising of "lack of information" and "unclear consequences" improved model fit over a flat model. Yet, the linear mixed models showed that introducing uncertain information had a stronger effect than value conflict when predicting "lack of information" and "unclear consequences" (i.e., the two facets of empirical uncertainty), but not when predicting "normative uncertainty" (see Supplementary Material). We took this as further evidence of the validity of our factors and the distinction between empirical and normative uncertainty.

General Discussion

While there has been much philosophical interest revolving normative uncertainty and its similarities and differences with empirical uncertainty, there is only limited psychological literature on the matter so far. In two studies, we showed that there is a measurable difference between being morally uncertain due to uncertain facts and being morally uncertain due to unclear or conflicting values. We introduce the Moral Uncertainty Scale which consists of factors that have been identified on a dataset with participant-generated scenarios and that have been validated on a set of researcher-generated scenarios that varied in scenario setting, the amount of given empirical information, and the amount of value conflict. The MUS showed strict measurement invariance across all these configurations and is thus expected to distinguish between both forms of moral uncertainty are predictive of overall moral certainty and that they are distinct from attitude clarity and attitude correctness. Lastly, we showed that the factors of the MUS are sensitive to respective manipulations of scenario content, which revealed interactions between uncertain information and conflicting values.

Both the idea of moral decision-making under empirical uncertainty and that of value conflict are not novel by themselves, since they have been considered in different psychological paradigms in a largely unconnected manner (Fleischhut et al., 2017; Hanselmann & Tanner, 2008; Kortenkamp & Moore, 2014; Mandel & Vartanian, 2008; Shaddy et al., 2021; Shou et al., 2020; Shou & Song, 2017). Yet, investigations of uncertainty in moral decision-making have considered neither the difference between these constructs nor that they can potentially occur simultaneously. We show that empirical and normative uncertainty can arise in the same situations and that both are relevant for explaining uncertainty in moral judgments. We thereby try to connect these largely separated approaches. By enabling researchers to identify whether participants are empirically and/or normatively uncertain in moral issues that the researchers might be interested in, we hope to foster novel research into the communalities and differences of these two constructs. Such research might address how the proposed distinction affects the resolution of uncertainty, attitude change, or the attitude-behavior gap.

Differentiating between empirical and normative uncertainty is further crucial for the study of folk metaethics. As Theisen (2023) recently showed, how lay people conceive of the nature of moral thought and language is associated with how they interpret and try to resolve normative uncertainty. However, one limitation of such studies is a potential confounding of normative uncertainty with empirical uncertainty (see also Bush & Moss, 2020). One way to address this concern is to find issues where the elicited uncertainty is assumed to be genuinely normative. Yet, this not only limits the researchers' choice in stimulus material, which might lead to biased results, but it has so far also not been possible to control whether the stimulus material elicits that kind of uncertainty the researchers are interested in. Being able to separately measure both constructs allows experimental philosophers to ask participants about moral issues that might elicit both empirical and normative uncertainty, which might increase internal and external validity of such studies and further inform us about how lay people conceive of and act under normative uncertainty.

Limitations

One potential limitation of the current research is that it largely focused on normative uncertainty arising from value conflict. Even though we included items referring to other aspects of normative uncertainty, such as doubt regarding the moral value of individual attributes, these did not emerge as a contingent factor in our studies. This might be explained by participants largely reporting scenarios that involved value conflict in Study 1. As we mentioned above, moral dilemmas might first come to mind when thinking about difficult moral cases, so that participants may not have thought of issues leading to other forms of normative uncertainty. Further, participants who were uncertain regarding the moral relevance of an action might have been hesitant to describe that action in Study 1 just because of their uncertainty. This would have reduced the scope of the generated scenarios and would explain why other facets of normative uncertainty did not emerge as clearly. It should also be noted that even though most of the normative uncertainty items refer to value conflict, this is not true for all items in our factor, so that the scale captures a broader concept, including aspects such as uncertainty regarding the application of one's moral views to a specific situation and uncertainty regarding ethical arguments.

Lastly, we could show that all uncertainty-related variables in Study 2 correlated highly with the overall degree of (un-)certainty. This seems to be a general problem when trying to assess different sources of uncertainty. As we have laid out, one cannot be empirically or normatively uncertain without also being generally uncertain about the respective moral judgment. Researchers might choose to control for the overall degree of uncertainty when using the MUS to differentiate between empirical and normative uncertainty.

Conclusion

As shown by the many reports from participants in our first study, being morally uncertain is a widespread phenomenon that can appear in private, professional, and political considerations. It is not always possible to fully resolve one's moral doubts. For moral psychology, it is therefore vital to understand what moral uncertainty actually constitutes. Building upon a growing tradition of philosophical literature, we advanced this understanding by showing that moral uncertainty can either refer to empirical or to normative aspects of the moral issue. We hope that this will foster further research into the antecedents and consequences of this distinction.

References

Alsaad, A. K. (2021). Ethical judgment, subjective norms, and ethical consumption: The moderating role of moral certainty. *Journal of Retailing and Consumer Services*, *59*, 102380. https://doi.org/10.1016/j.jretconser.2020.102380

Alsaad, A. K., Elrehail, H., & Saif-Alyousfi, A. Y. H. (2021). The interaction among religiosity, moral intensity and moral certainty in predicting ethical consumption: A study of Muslim consumers. *International Journal of Consumer Studies*, *46*, 406-418. https://doi.org/10.1111/ijcs.12688

- Alvarez, R. M., & Brehm, J. (1995). American ambivalence towards abortion policy: Development of a Heteroskedastic Probit Model of competing values. *American Journal of Political Science*, *39*(4), 1055–1082.
- Alvarez, R. M., & Brehm, J. (1997). Are Americans ambivalent towards racial policies? *American Journal of Political Science*, *41*(2), 345–374.
- Andrejević, M., Feuerriegel, D., Turner, W., Laham, S., & Bode, S. (2020). Moral judgements of fairness-related actions are flexibly updated to account for contextual information. *Scientific Reports*, 10(1), 17828. https://doi.org/10.1038/s41598-020-74975-0
- Bartlett, M. S. (1950). Tests of significance in factor analysis. *British Journal of Statistical Psychology*, *3*(2), 77–85. https://doi.org/10.1111/j.2044-8317.1950.tb00285.x
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using Ime4. *Journal of Statistical Software*, *67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01

Bauman, C. W., McGraw, A. P., Bartels, D. M., & Warren, C. (2014). Revisiting external validity:
Concerns about trolley problems and other sacrificial dilemmas in moral psychology.
Social and Personality Psychology Compass, 8(9), 536–554.
https://doi.org/10.1111/spc3.12131

Białek, M., & Neys, W. de (2017). Dual processes and moral conflict: Evidence for deontological reasoners' intuitive utilitarian sensitivity. *Judgment and Decision Making*, 12(2), 148– 167. https://doi.org/10.1017/S1930297500005696

- Biek, M., Wood, W., & Chaiken, S. (1996). Working knowledge, cognitive processing, and attitudes: On the determinants of bias. *Personality & Social Psychology Bulletin*, 22(6), 547–556. https://doi.org/10.1177/0146167296226001
- Bourget, D., & Chalmers, D. J. (2023). Philosophers on philosophy: The 2020 PhilPapers Survey. *Philosopher's Imprint*, *23*(1). https://doi.org/10.3998/phimp.2109
- Bush, L. S., & Moss, D. (2020). Misunderstanding metaethics: Difficulties measuring folk objectivism and relativism. *Diametros*, 17(64), 6–21. https://doi.org/10.33392/diam.1495
- Chalmers, R. P. (2012). mirt: A multidimensional Item Response Theory package for the R environment. *Journal of Statistical Software*, *48*(6), 1–29. https://doi.org/10.18637/jss.v048.i06
- Clarkson, E., & Jasper, J. D. (2022). Individual differences in moral judgment predict attitudes towards mandatory vaccinations. *Personality and Individual Differences, 186*, 111391. https://doi.org/10.1016/j.paid.2021.111391
- Coelho, G. L. d. H., Hanel, P. H. P., & Wolf, L. J. (2020). The very efficient assessment of Need for Cognition: Developing a six-item version. *Assessment*, *27*(8), 1870–1885. https://doi.org/10.1177/1073191118793208

Costa-Gomes, M., & Schoenegger, P. (2023). Charitable giving under normative uncertainty: Experimental evidence on the behavioral impact of normative expert advice. *SSRN Electronic Journal.* Advance online publication. https://doi.org/10.2139/ssrn.4389737 Crutchfield, P., Scheall, S., Rzeszutek, M. J., Brown, H. D., & Cardoso Sao Mateus, C. (2023).

Ignorance and moral judgment: Testing the logical priority of the epistemic.

Consciousness and Cognition, 108, 103472.

https://doi.org/10.1016/j.concog.2023.103472

Dietrich, F., Fleurbaey, M., & Jabarian, B. (2019). *The Normative Uncertainty Survey.* 2nd *Workshop on Global Priorities Research*, Oxford, UK.

Ellemers, N., van der Toorn, J., Paunov, Y., & van Leeuwen, T. (2019). The psychology of morality:
A review and analysis of empirical studies published from 1940 through 2017.
Personality and Social Psychology Review, 23(4), 332–366.
https://doi.org/10.1177/10888683188117

- Eriksson, J., & Francén Olinder, R. (2016). Non-cognitivism and the classification account of moral uncertainty. *Australasian Journal of Philosophy*, 94(4), 719–735. https://doi.org/10.1080/00048402.2016.1149199
- Feinberg, M., Kovacheff, C., Teper, R., & Inbar, Y. (2019). Understanding the process of moralization: How eating meat becomes a moral issue. *Journal of Personality and Social Psychology*, 117(1), 50–72. https://doi.org/10.1037/pspa0000149
- Fernández, L. (2021). Images that liberate: Moral shock and strategic visual communication in animal liberation activism. *Journal of Communication Inquiry*, 45(2), 138–158. https://doi.org/10.1177/0196859920932881
- Fleischhut, N., Meder, B., & Gigerenzer, G. (2017). Moral hindsight. *Experimental Psychology*, 64(2), 110–123. https://doi.org/10.1027/1618-3169/a000353

Gerpott, F. H., Balliet, D., Columbus, S., Molho, C., & Vries, R. E. de (2018). How do people think about interdependence? A multidimensional model of subjective outcome interdependence. *Journal of Personality and Social Psychology*, *115*(4), 716–742. https://doi.org/10.1037/pspp0000166 Greene, J. D., Cushman, F. A., Stewart, L. E., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2009). Pushing moral buttons: The interaction between personal force and intention in moral judgment. *Cognition*, *111*(3), 364–371. https://doi.org/10.1016/j.cognition.2009.02.001

- Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, 293(5537), 2105– 2108. https://doi.org/10.1126/science.1062191
- Hanselmann, M., & Tanner, C. (2008). Taboos and conflicts in decision making: Sacred values, decision difficulty, and emotions. *Judgment and Decision Making*, *3*(1), 51–63. https://doi.org/10.5167/uzh-5948

Hansson, S. O. (2013). The ethics of risk. Palgrave Macmillan UK.

https://doi.org/10.1057/9781137333650

- Hazlett, C., & Wainstein, L. (2022). Understanding, choosing, and unifying multilevel and fixed effect approaches. *Political Analysis*, *30*(1), 46–65. https://doi.org/10.1017/pan.2020.41
- Hendrickson, A. E., & White, P. O. (1964). Promax: A quick method for rotation to oblique simple structure. *British Journal of Statistical Psychology*, *17*(1), 65–70.

https://doi.org/10.1111/j.2044-8317.1964.tb00244.x

- Hicks, A. (2018). Moral uncertainty and value comparison. In R. Shafer-Landau (Ed.), Oxford studies in metaethics (pp. 161–183). Oxford University Press. https://doi.org/10.1093/oso/9780198823841.003.0008
- Horn, J. L. (1965). A rationale and test for the number of factors in factor analysis. *Psychometrika*, *30*, 179–185. https://doi.org/10.1007/BF02289447
- Jabarian, B. (2020, May 30). *Moral uncertainty, social values and economic preferences* [Paper presentation]. Princeton University, Princeton, NJ, USA.
- Kaiser, H. F. (1974). An index of factorial simplicity. *Psychometrika*, *39*(1), 31–36. https://doi.org/10.1007/BF02291575

- Keele, L., & Park, D. K. (2006). *Difficult choices: An evaluation of heterogeneous choice models*. https://www3.nd.edu/~rwilliam/oglm/ljk-021706.pdf
- Kodapanakkal, R. I., Brandt, M. J., Kogler, C., & Beest, I. (2022). Moral relevance varies due to inter-individual and intra-individual differences across big data technology domains. *European Journal of Social Psychology*, *52*(1), 46–70. https://doi.org/10.1002/ejsp.2814
- Kolenikov, S., & Bollen, K. A. (2012). Testing negative error variances. *Sociological Methods & Research*, *41*(1), 124–167. https://doi.org/10.1177/0049124112442138
- Kortenkamp, K. V., & Moore, C. F. (2014). Ethics under uncertainty: The morality and appropriateness of utilitarianism when outcomes are uncertain. *American Journal of Psychology*, *127*(3), 367. https://doi.org/10.5406/amerjpsyc.127.3.0367
- Leiner, D. J. (2019). Too fast, too straight, too weird: Non-reactive indicators for meaningless data in internet surveys. *Survey Research Methods*, *13*(3), 229–248. https://doi.org/10.18148/srm/2019.v13i3.7403
- Leiner, D. J. (2022). *SoSci Survey* (Version 3.2.55) [Computer software]. https://www.soscisurvey.de

https://doi.org/10.2307/2290157

Little, R. J. A. (1988). A test of missing completely at random for multivariate data with missing values. *Journal of the American Statistical Association*, *83*(404), 1198–1202.

Lockhart, T. (2000). *Moral uncertainty and its consequences*. Oxford University Press.

- Lowth, M. (2020). Can mandatory vaccination be ethically justified? A deontological perspective. *King's Student Law Review*, *11*(2), 26–48.
- Luttrell, A., Petty, R. E., & Briñol, P. (2020). The interactive effects of ambivalence and certainty on political opinion stability. *Journal of Social and Political Psychology*, 8(2), 525–541. https://doi.org/10.5964/jspp.v8i2.1247
- MacAskill, W. (2016). Normative uncertainty as a voting problem. *Mind*, *125*(500), 967–1004. https://doi.org/10.1093/mind/fzv169

MacAskill, W., Bykvist, K., & Ord, T. (Eds.). (2020). Moral uncertainty. Oxford University Press.

- Makins, N. (2021). Attitudinal ambivalence: Moral uncertainty for non-cognitivists. *Australasian Journal of Philosophy*, 1–15. https://doi.org/10.1080/00048402.2021.1908380
- Mandel, D. R., & Vartanian, O. (2008). Taboo or tragic: effect of tradeoff type on moral choice, conflict, and confidence. *Mind & Society*, *7*(2), 215–226. https://doi.org/10.1007/s11299-007-0037-3
- Mata, A. (2019). Social metacognition in moral judgment: Decisional conflict promotes perspective taking. *Journal of Personality and Social Psychology*, *117*(6), 1061–1082. https://doi.org/10.1037/pspa0000170
- McNeish, D., Stapleton, L. M., & Silverman, R. D. (2017). On the unnecessary ubiquity of hierarchical linear modeling. *Psychological Methods*, 22(1), 114–140. https://doi.org/10.1037/met0000078
- Meder, B., Fleischhut, N., Krumnau, N.-C., & Waldmann, M. R. (2019). How should autonomous cars drive? A preference for defaults in moral judgments under risk and uncertainty. *Risk Analysis*, *39*(2), 295–314. https://doi.org/10.1111/risa.13178
- Moller, D. (2011). Abortion and moral risk. *Philosophy*, 86(3), 425–443.

https://doi.org/10.1017/S0031819111000222

Monroe, A. E., & Malle, B. F. (2019). People systematically update moral judgments of blame. Journal of Personality and Social Psychology, 116(2), 215–236. https://doi.org/10.1037/pspa0000137

Mundfrom, D. J., Shaw, D. G., & Ke, T. L. (2005). Minimum sample size recommendations for conducting factor analyses. *International Journal of Testing*, 5(2), 159–168. https://doi.org/10.1207/s15327574ijt0502_4

Nissan-Rozen, I. (2015). Against moral hedging. *Economics and Philosophy*, *31*(3), 349–369. https://doi.org/10.1017/S0266267115000206

- O'Connor, A. M. (1995). Validation of a Decisional Conflict Scale. *Medical Decision Making*, 15, 25–30. https://doi.org/10.1177/0272989x9501500105
- Petrocelli, J. V., Tormala, Z. L., & Rucker, D. D. (2007). Unpacking attitude certainty: Attitude clarity and attitude correctness. *Journal of Personality and Social Psychology*, *92*(1), 30–41. https://doi.org/10.1037/0022-3514.92.1.30
- Philipp-Muller, A. Z., Wallace, L. E., & Wegener, D. T. (2020). Where does moral conviction fit? A factor analytic approach examining antecedents to attitude strength. *Journal of Experimental Social Psychology*, *86*, 103900. https://doi.org/10.1016/j.jesp.2019.103900
- Poliquin, C. W. (2010). Protected values, range effects, guilt, and tradeoff difficulty in moral decision making. *Penn Journal of Philosophy, Politics & Economics*, 5(1), Article 5. https://repository.upenn.edu/spice/vol5/iss1/5
- R Core Team. (2020). *R: A language and environment for statistical computing* (Version 4.0.3) [Computer software]. Vienna, Austria. https://www.R-project.org/
- Reich, T., & Wheeler, S. C. (2016). The good and bad of ambivalence: Desiring ambivalence under outcome uncertainty. *Journal of Personality and Social Psychology*, *110*(4), 493–508. https://doi.org/10.1037/pspa0000047
- Reise, S. P., Ventura, J., Nuechterlein, K. H., & Kim, K. H. (2005). An illustration of multilevel factor analysis. *Journal of Personality Assessment*, 84(2), 126–136. https://doi.org/10.1207/s15327752jpa8402_02
- Revelle, W. (2021). *psych: Procedures for psychological, psychometric, and personality research.* Northwestern University. https://cran.r-project.org/package=psych
- Rhee, J. J., Schein, C., & Bastian, B. (2019). The what, how, and why of moralization: A review of current definitions, methods, and evidence in moralization research. *Social and Personality Psychology Compass*, 13(12). https://doi.org/10.1111/spc3.12511

Robinson, P., & Steele, K. (2022). Moral uncertainty, noncognitivism, and the multi-objective story. *Noûs*, Article nous.12435. Advance online publication.

https://doi.org/10.1111/nous.12435

Rönkkö, M., & Cho, E. (2022). An updated guideline for assessing discriminant validity. *Organizational Research Methods*, 25(1), 6–14. https://doi.org/10.1177/1094428120968614

Rosenthal, C. (2019). Ethics for fallible people [Dissertation]. New York University, New York, NY.

Rosseel, Y. (2012). lavaan : An R package for structural equation modeling. *Journal of Statistical Software*, 48(2). https://doi.org/10.18637/jss.v048.i02

Rowland, R. (2021). Moral disagreement. Routledge.

Rudnev, M., Lytkina, E., Davidov, E., Schmidt, P., & Zick, A. (2018). Testing measurement invariance for a second-order factor. A cross-national test of the Alienation Scale. *methods, data, analyses, 12*(1), 47-76. https://doi.org/10.12758/mda.2017.11

- Samejima, F. (1969). Estimation of latent ability using a response pattern of graded scores. *Psychometrika*, *34*(S1), 1–97. https://doi.org/10.1007/bf03372160
- Savulescu, J. (2021). Good reasons to vaccinate: Mandatory or payment for risk? *Journal of Medical Ethics*, 47(2), 78–85. https://doi.org/10.1136/medethics-2020-106821
- Schein, C. (2020). The importance of context in moral judgments. Perspectives on Psychological Science: A Journal of the Association for Psychological Science, 15(2), 207–215. https://doi.org/10.1177/1745691620904083

Shaddy, F., Fishbach, A., & Simonson, I. (2021). Trade-offs in choice. *Annual Review of Psychology*, *72*, 181–206. https://doi.org/10.1146/annurev-psych-072420-125709

Shou, Y., Olney, J., Smithson, M., & Song, F. (2020). Impact of uncertainty and ambiguous outcome phrasing on moral decision-making. *PloS One*, *15*(5), e0233127. https://doi.org/10.1371/journal.pone.0233127

- Shou, Y., & Song, F. (2017). Decisions in moral dilemmas: The influence of subjective beliefs in outcome probabilities. *Judgment and Decision Making*, *12*(5), 481–490.
- Skitka, L. J. (2010). The psychology of moral conviction. *Social and Personality Psychology Compass*, 4(4), 267–281. https://doi.org/10.1111/j.1751-9004.2010.00254.x
- Skovgaard-Olsen, N., & Klauer, K. C. (2024). Norm conflicts and morality: The CNIS conflict model of moral decision-making [Unpublished manuscript].
- Tamborini, R., Prabhu, S., Lewis, R. J., Grizzard, M., & Eden, A. (2018). The influence of media exposure on the accessibility of moral intuitions and associated affect. *Journal of Media Psychology*, 30(2), 79–90. https://doi.org/10.1027/1864-1105/a000183
- Tarsney, C. (2018). Moral uncertainty for deontologists. *Ethical Theory and Moral Practice*, *21*(3), 505–520. https://doi.org/10.1007/s10677-018-9924-4
- Tarsney, C. (2021). Normative externalism by Brian Weatherson. *Mind*, *130*(519), 1018–1028. https://doi.org/10.1093/mind/fzaa018
- Theisen, M. (2023). Metaethical intuitions in lay concepts of normative uncertainty. *Philosophical Psychology.* Advance online publication. https://doi.org/10.1080/09515089.2023.2174422
- Tormala, Z. L., & Rucker, D. D. (2018). Attitude certainty: Antecedents, consequences, and new directions. *Consumer Psychology Review*, 1(1), 72–89.

https://doi.org/10.1002/arcp.1004

van Buuren, S., & Groothuis-Oudshoorn, K. (2011). mice: Multivariate imputation by chained equations in R. *Journal of Statistical Software*, *45*(3).

https://doi.org/10.18637/jss.v045.i03

van Harreveld, F., van der Pligt, J., & Liver, Y. N. de (2009). The agony of ambivalence and ways to resolve it: Introducing the MAID model. *Personality and Social Psychology Review*, *13*(1), 45–61. https://doi.org/10.1177/1088868308324518
Vega, S., Mata, A., Ferreira, M. B., & Vaz, A. R. (2020). Metacognition in moral decisions:
 judgment extremity and feeling of rightness in moral intuitions. *Thinking & Reasoning*,
 1–18. https://doi.org/10.1080/13546783.2020.1741448

Watkins, M. W. (2018). Exploratory factor analysis: A guide to best practice. *Journal of Black Psychology*, 44(3), 219–246. https://doi.org/10.1177/0095798418771807

Weatherson, B. (2019). Normative externalism. Oxford University Press.

Wood, W., Nancy Rhodes, & Biek, M. (1995). Working knowledge and attitude strength: An information-processing analysis. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 189–202). Psychology Press.

Young, L., & Saxe, R. (2009). An FMRI investigation of spontaneous mental state inference for moral judgment. *Journal of Cognitive Neuroscience*, *21*(7), 1396–1405. https://doi.org/10.1162/jocn.2009.21137

Overview of Key Concepts

Concept	Definition
Moral uncertainty	Doubt or lack of confidence or clarity regarding one's moral judgment of a certain act.
Empirical uncertainty	Moral uncertainty that arises from ignorance of empirical facts, such as the empirical circumstances or potential consequences of an act.
Normative uncertainty	Moral uncertainty that arises from a lack of confidence regarding the relevant moral standards, their relative weight, and/or their application to the concrete situation.

Descriptive Statistics, Pattern Coefficients, and Communality for Items Included in Final EFA in

Study 1

					Factors		
ltem	Text	M (SD)	Skew (Kurt)	Info	Cons	Norm	h²
EU01	It is difficult to judge the action because I do not know what doing or not doing the action would lead to.	4.02 (2.00)	-0.07 (-1.34)	0.03	0.80	-0.02	0.66
EU02	I would need to know more about the consequences of doing or not doing the action in order to judge its appropriateness.	3.97 (1.92)	-0.05 (-1.28)	0.21	0.59	0.01	0.57
EU03	One would need more information to be able to judge the situation clearly.	3.57 (1.99)	0.21 (-1.29)	0.72	0.22	-0.08	0.70
EU04	It is difficult to make a moral decision because I would not know all the relevant facts in the situation.	3.48 (1.87)	0.19 (-1.21)	0.77	0.13	-0.02	0.71
EU05	I lack important information to be able to judge the action.	3.17 (1.84)	0.50 (-0.93)	0.82	0.07	0.02	0.77
EU06	It would be easier to make a moral decision if I knew more precisely what effects doing or not doing the action would have.	4.72 (1.95)	-0.58 (-0.86)	0.05	0.66	0.00	0.48
EU07	I would need more facts to make a decision in this situation.	3.27 (1.91)	0.47 (-0.99)	0.80	0.05	0.01	0.71
EU08	It is difficult to make a moral decision because I don't know what has led to the situation.	2.67 (1.96)	0.94 (-0.51)	0.69	-0.33	0.14	0.36

					Factors		
Item	Text	M (SD)	Skew (Kurt)	Info	Cons	Norm	h²
EU09	Since I do not know everything about the situation, it is difficult to judge the action from a moral point of view.	3.24 (1.71)	0.35 (-0.92)	0.84	0.03	0.01	0.75
EU10	The moral evaluation of the action is difficult because some facts are unclear.	3.52 (1.93)	0.22 (-1.27)	0.62	0.27	-0.03	0.64
EU11	I have all the relevant facts to judge the appropriateness of the action. (reversed)	3.77 (1.82)	0.17 (-1.07)	0.49	0.20	-0.04	0.38
EU12	It is foreseeable what morally significant consequences it would have if one does or does not do the action. (reversed)	3.69 (1.79)	0.26 (-0.93)	-0.06	0.45	-0.05	0.16
EU13	The moral evaluation of the action is so difficult because some consequences are unclear.	4.34 (2.10)	-0.31 (-1.32)	-0.04	0.79	-0.01	0.59
EU14	Since I do not know in advance the exact consequences of doing or not doing the action, it is difficult to judge the action from a moral point of view.	3.66 (1.86)	0.12 (-1.21)	0.11	0.67	0.03	0.59
EU15	I am not sure what exactly the consequences of doing or not doing the action would be.	4.00 (1.89)	-0.13 (-1.29)	-0.08	0.73	-0.02	0.45
EU16	I am not able to predict the morally significant consequences of doing or not doing the action.	3.72 (1.86)	0.23 (-1.10)	-0.02	0.65	0.11	0.48
EU17	It is difficult to judge the action because the background of the situation is not clear.	2.98 (1.86)	0.64 (-0.85)	0.76	-0.02	-0.09	0.51

					Factors		
Item	Text	M (SD)	Skew (Kurt)	Info	Cons	Norm	h²
NU01	Even if I knew everything about the situation, there would be significant doubts as to how to evaluate it morally.	3.46 (1.92)	0.19 (-1.24)	-0.02	-0.07	0.67	0.41
NU02	I do not know which of the arguments for or against the appropriateness of the action should be the decisive one.	3.84 (1.96)	0.01 (-1.30)	0.04	0.00	0.65	0.44
NU04	I am unsure how much I should consider the various aspects of this situation in my judgment.	4.36 (1.76)	-0.38 (-0.95)	0.05	0.08	0.60	0.45
NU05	I have different moral views that lead to different evaluations of this action.	4.50 (1.89)	-0.38 (-1.03)	-0.04	-0.06	0.77	0.54
NU06	I am unsure how to evaluate the consequences of doing or not doing the action from a moral point of view.	4.10 (1.78)	-0.14 (-1.17)	0.04	0.48	0.14	0.34
NU07	Depending on the perspective from which I think about the situation, I arrive at different evaluations of the action.	4.97 (1.78)	-0.60 (-0.72)	0.08	-0.08	0.71	0.51
NU08	I am unsure which moral view I should adhere to when judging the action.	3.94 (1.88)	-0.06 (-1.17)	-0.01	0.07	0.76	0.62
NU10	Conflicting arguments come to my mind when judging this action.	5.08 (1.82)	-0.74 (-0.55)	-0.10	0.13	0.74	0.56
NU11	It is possible that I am wrong in the moral assessment of the situation.	4.52 (1.75)	-0.39 (-0.79)	0.08	-0.02	0.48	0.26
NU13	My moral views do not fit together in this situation.	3.98 (1.85)	-0.02 (-1.15)	-0.09	0.01	0.59	0.31
NU14	I am unsure how to apply my moral concepts to this situation.	3.86 (1.84)	-0.11 (-1.13)	0.03	0.09	0.68	0.54

Note. Bold printed items were selected for the final scale. Skew: skewness, Kurt: kurtosis, *h*²: communality, Info: "lack of information", Cons: "unclear consequences", Norm: "normative uncertainty".

Variable	Overall	University	Prolific	Group comparison
	(N - 102)	(n - 1.16)	(n - 256)	
	(10 - 402)	(11 - 140)	(11 - 250)	
Age	27.9	22.2	31.0	<i>t</i> (383.7 ^a) = -11.7, <i>p</i> < .001, <i>d</i> = 1.02
0	(9.61)	(4.68)	(10.2)	
	(5.01)	(4.00)	(10.2)	
Gender				$\chi^2 = 79.2, \rho^{\rm b} < .001, V(1) = .442$
Male	160	16	144	
indic	(20.99/)	(11.0%)	(EG 20/)	
	(59.6%)	(11.0%)	(50.5%)	
Female	231	124	107	
		(04.00/)	(11 00/)	
	(57.5%)	(84.9%)	(41.8%)	
Non-binary/	5	1	4	
Diverse	(1.2%)	(0.7%)	(1.6%)	
Diverse	(1.270)	(0.770)	(1.070)	
Completion	26.6	29.3	25.0	t(400) = 4.69, p < .001, d = 0.49
time (minutes)	(0 15)	(9 68)	(8 16)	
time (minutes)	(5.15)	(5.08)	(0.40)	

Description of Final Sample in Study 2

Note. Mean (standard deviation) is given for age and completion time, other variables are

described in absolute numbers (percentages).

^a Welch-corrected for unequal variances. ^b *p*-value based on Monte Carlo simulation due to small

cell sizes.

Variable	Mean	SD	Skewness	Kurtosis	Cronbach's α
Empirical uncertainty	3.34	0.14	0.42	-0.98	.972
Lack of information	3.39	0.14	0.36	-1.87	.965
Unclear consequences	3.27	0.11	0.05	-2.00	.945
Normative uncertainty	3.17	0.23	0.3	-1.84	.950
Overall certainty	4.36	0.19	0.35	-2.02	.940
Clarity	5.36	0.10	0.15	-2.27	.955
Correctness	4.41	0.17	0.12	-2.33	.926

Descriptive Statistics for All Variables in Study 2

Note. Unweighted means.

Model	CFI robust	TLI robust	RMSEA robust	RMSEA _{robust} 90% Cl	SRMR	
	Scenario	Setting: Oil v	vs. Company			
Overall model	0.975	0.970	.069	[.065; .072]	0.026	
"Oil" model	0.971	0.966	.069	[.064; .074]	0.029	
"Company" model	0.976	0.972	.068	[.062; .074]	0.027	
Configural model	0.974	0.969	.068	[.065; .072]	0.028	
Metric model (1 st order)	0.974	0.971	.067	[.063; .071]	0.030	
Metric model (1 st & 2 nd order)	0.974	0.971	.067	[.063; .070]	0.030	
Scalar model (1 st order)	0.974	0.972	.065	[.061; .069]	0.030	
Scalar model (1 st & 2 nd order)	0.973	0.972	.065	[.061; .069]	0.030	
Strict model	0.971	0.972	.066	[.062; .069]	0.031	
Empirical Uncertainty Manipulation: Low vs. High						
Overall model	0.975	0.970	.069	[.065; .072]	0.026	
"Low" model ^a	0.977	0.973	.070	[.064; .076]	0.024	
"High" model	0.971	0.965	.067	[.061; .072]	0.031	
Configural model ^a	0.975	0.970	.068	[.064; .072]	0.027	
Metric model (1 st order) ^a	0.974	0.971	.067	[.063; .070]	0.030	
Metric model (1 st & 2 nd order) ^a	0.974	0.971	.067	[.063; .070]	0.030	
Scalar model (1 st order) ^a	0.972	0.971	.067	[.063; .070]	0.030	
Scalar model (1 st & 2 nd order)	0.972	0.971	.067	[.063; .070]	0.030	
Strict model ^a	0.967	0.968	.070	[.066; .073]	0.031	
Noi	rmative Unce	rtainty Manip	ulation: Low vs	. High		
Overall model	0.975	0.970	.069	[.065; .072]	0.026	
"Low" model	0.977	0.973	.066	[.060; .071]	0.024	
"High" model	0.973	0.969	.069	[.063; .074]	0.027	
Configural model	0.975	0.971	.067	[.063; .071]	0.025	

Model Fits for Different Levels of Measurement Invariance in Study 2

Model	<i>CFI</i> _{robust}	<i>TLI</i> robust	<i>RMSEA</i> robust	RMSEA _{robust} 90% Cl	SRMR
Metric model (1 st order)	0.975	0.972	.066	[.062; .070]	0.031
Metric model (1 st & 2 nd order)	0.975	0.972	.066	[.062; .069]	0.031
Scalar model (1 st order)	0.972	0.971	.067	[.063; .071]	0.035
Scalar model (1 st & 2 nd order)	0.972	0.971	.067	[.063; .071]	0.035
Strict model	0.972	0.972	.065	[.062; .069]	0.035

Note. Invariance given if difference in *CFI*_{robust} between parsimonious and less parsimonious model <= 0.01. "Oil" and "company" refer to models where only "oil" or only "company" scenarios were considered. "Low" and "high" refer to models where only scenarios with "low" or "high" empirical/normative uncertainty manipulation have been considered.

^a Yielded negative variance estimates that were not significantly different from zero.

Fixed and Random Effects for Predicting Overall Moral Certainty From Empirical Uncertainty,

		95%		
Fixed effects	b	Lower limit	Upper limit	$R^2_{partial}$
Intercept	1.144	0.799	1.490	-
Clarity	0.322	0.258	0.386	3.4%
Correctness	0.222	0.124	0.319	0.7%
Empirical uncertainty	-0.233	-0.267	-0.198	6.4%
Clarity × Correctness	0.052	0.036	0.068	1.3%
Random effects	Variance	Correlation		
Participant		Intercept		
Random intercept	0.091			
Random slope: Empirical uncertainty	0.009	.09		
Residual	0.781			

Clarity, and Correctness in Study 2

Note. N = 402, 3216 observations, $R^{2}_{fixed} = 74.75\%$, $R^{2}_{total} = 80.52\%$.

Fixed and Random Effects for Predicting Overall Moral Certainty From Normative Uncertainty,

		95%		
Fixed effects	b	Lower limit	Upper limit	$R^2_{partial}$
Intercept	2.430	2.057	2.805	-
Clarity	0.270	0.207	0.333	2.7%
Correctness	0.243	0.129	0.320	0.8%
Normative uncertainty	-0.441	-0.480	-0.402	15.6%
Clarity × Correctness	0.036	0.021	0.052	0.7%
Random effects	Variance			
Participant				
Random intercept	0.068			
Random slope:	0.010			
Normative uncertainty				
Residual	0.720			

Clarity, and Correctness in Study 2

Note. N = 402, 3216 observations, $R^{2}_{fixed} = 79.19\%$, $R^{2}_{total} = 81.24\%$.

Fixed and Random Effects for Predicting Overall Moral Certainty From MUS Factors, Clarity, and

Fixed effects	b	Lower limit	Upper limit	$R^2_{partial}$
Intercept	2.627	2.272	2.984	-
Clarity	0.299	0.238	0.359	3.3%
Correctness	0.246	0.154	0.338	0.9%
Empirical uncertainty	-0.098	-0.133	-0.063	1.1%
Normative uncertainty	-0.403	-0.445	-0.362	11.0%
Clarity × Correctness	0.027	0.011	0.042	0.4%
Random effects	Variance	Correlation		
Participant		Intercept		
Random intercept	0.020			
Random slope: Empirical uncertainty	0.007	.92		
Residual	0.711			

Correctness in Study 2

Note. N = 402, 3216 observations, $R^{2}_{fixed} = 77.13\%$, $R^{2}_{total} = 82.00\%$.

Latent Correlations Between MUS Factors and Clarity/Correctness

Variable	Empirical uncertainty	Normative uncertainty	Clarity	Correctness
Empirical uncertainty	-	.781	583	669
Normative uncertainty	.493	-	714	766
Clarity	095	225	-	.766
Correctness	220	272	.381	-

Note. Upper triangle: full correlations; lower triangle: partial correlations controlling for overall

moral certainty.

Fixed and Random Effects for Predicting Empirical and Normative Uncertainty Scores From

Empirical uncertainty score ^a	npirical uncertainty score ^a 95% Cl			
Fixed effects	b	Lower limit	Upper limit	R² _{partial}
Intercept	2.232	2.118	2.346	-
Empirical uncertainty manipulation	1.318	1.210	1.427	10.4%
Normative uncertainty manipulation	0.841	0.723	0.959	4.5%
Setting (Oil)	0.690	0.595	0.785	5.9%
Emp. × norm. uncertainty manipulation	-1.308	-1.461	-1.155	5.4%
Random effects	Variance	Correl	ations	
Participant		Intercept	Norm. unc. manipulation	
Random intercept	0.586			
Random slope: Norm. unc. manipulation	0.219	.22		
Random slope: Setting (Oil)	0.334	16	77	
Residual	1.231			
Normative uncertainty score	e ^b	95%	% CI	
Fixed effects	b	Lower limit	Upper limit	R^2_{partial}
Intercept	1.813	1.710	1.916	-
Empirical uncertainty manipulation	1.346	1.247	1.444	12.7%
Normative uncertainty manipulation	1.572	1.454	1.691	16.5%
Setting (Oil)	0.625	0.536	0.715	5.9%
Emp. × norm. uncertainty manipulation	-1.643	-1.782	-1.504	9.7%

Manipulations of Empirical and Normative Uncertainty in Study 2

114 PHILOSOPHICAL AND LAY CONCEPTS OF MORAL UNCERTAINTY

Random effects	Variance	Corre	elations	
Participant		Intercept	Norm. unc. manipulation	
Random intercept	0.485			
Random slope: Norm. unc. manipulation	0.464	.00		
Random slope: Setting (Oil)	0.332	31	60	
Residual	1.007			

Note. *N* = 402, 3216 observations.

 ${}^{a}R^{2}_{fixed}$ = 14.75%, R^{2}_{total} = 47.40%. ${}^{b}R^{2}_{fixed}$ = 22.41%, R^{2}_{total} = 52.62%.

Fixed and Random Effects for Predicting Overall Moral Certainty From Manipulations of Empirical

		95% CI		
Fixed effects	b	Lower limit	Upper limit	R^2_{partial}
Intercept	6.316	6.192	6.440	-
Empirical uncertainty manipulation	-2.052	-2.189	-1.915	16.8%
Normative uncertainty manipulation	-1.887	-2.043	-1.730	14.6%
Setting (Oil)	-1.103	-1.228	-0.977	10.4%
Emp. × norm. uncertainty manipulation	2.408	2.214	2.601	12.2%
Random effects	Variance	Correlations		
Participant		Intercept	Norm. unc. manipulation	
Random intercept	0.385			
Random slope: Norm. unc. manipulation	0.603	.21		
Randol slope: Setting (Oil)	0.667	40	59	
Residual	1.967			

and Normative Uncertainty in Study 2

Note. N = 402, 3216 observations, $R^{2}_{fixed} = 26.33\%$, $R^{2}_{total} = 46.22\%$.

Fixed and Random Effects for Predicting Empirical and Normative Uncertainty Scores From

Manipulations of Empirical and Normative Uncertainty While Controlling for Overall Moral

Certainty in Study 2

Empirical uncertainty score ^a		95% CI		
Fixed effects	b	Lower limit	Upper limit	$R^2_{partial}$
Intercept	5.287	5.121	5.453	-
Empirical uncertainty manipulation	0.325	0.226	0.424	0.9%
Normative uncertainty manipulation	-0.072	-0.173	0.029	0%
Setting (Oil)	0.156	0.080	0.233	0.5%
Overall certainty	-0.484	-0.506	-0.462	34.4%
Emp. × norm. uncertainty manipulation	-0.143	-0.279	-0.008	0.1%

Random effects	Variance	Corre	elations	
Participant		Intercept	Norm. unc. manipulation	
Random intercept	0.314			
Random slope: Norm. unc. manipulation	0.072	.44		
Random slope: Setting (Oil)	0.140	02	37	
Residual	0.812			

Normative uncertainty score ^b		95%	95% CI		
Fixed effects	b	Lower limit	Upper limit	$R^2_{partial}$	
Intercept	5.216	5.082	5.351	-	
Empirical uncertainty manipulation	0.240	0.161	0.318	0.8%	
Normative uncertainty manipulation	0.556	0.471	0.641	4.3%	
Setting (Oil)	0.031	-0.030	0.093	0%	
Overall certainty	-0.539	-0.557	-0.521	51.6%	
Emp. × norm. uncertainty manipulation	-0.346	-0.453	-0.238	0.9%	

Random effects	Variance	Corre	lations	
Participant		Intercept	Norm. unc. manipulation	
Random intercept	0.241			
Random slope: Norm. unc. manipulation	0.129	20		
Random slope: Setting (Oil)	0.099	21	43	
Residual	0.505			

Note. *N* = 402, 3216 observations.

 ${}^{a}R^{2}_{fixed}$ = 43.31%, R^{2}_{total} = 63.92%. ${}^{b}R^{2}_{fixed}$ = 62.09%, R^{2}_{total} = 75.06%.



Path Diagram With Standardized Estimates for the Full Model of Observed Variables in Study 2



Note. Intercepts omitted from plot for readability. Variances for EU, NU, Clarity, Certainty, and Correctness fixed to 1. EU: "empirical uncertainty"; NU: "normative uncertainty"; Info: "lack of information"; Cons: "unclear consequences".

Figure 2

Interaction Plots for Predicting Empirical and Normative Uncertainty Scores in Study 2



Note. A) Mean values of the empirical uncertainty score as a function of the empirical uncertainty manipulation (x-axis) and the normative uncertainty manipulation (color/shape).
B) Mean values of the normative uncertainty score as a function of the normative uncertainty manipulation (x-axis) and the empirical uncertainty manipulation (color/shape).

Figure 3

Interaction Plots for Predicting Empirical and Normative Uncertainty Scores After Controlling for



Overall Moral Certainty in Study 2

Note. A) Mean values of the empirical uncertainty score as a function of the empirical uncertainty manipulation (x-axis) and the normative uncertainty manipulation (color/shape).
B) Mean values of the normative uncertainty score as a function of the normative uncertainty manipulation (x-axis) and the empirical uncertainty manipulation (color/shape).

Appendix A: IRT Analysis of Items in Study 1

After the EFA, each factor consisted of at least nine items. In the following steps, we tried to reduce the number of items for each factor without negatively affecting the factor characteristics. We started by investigating the item intercorrelations. In the "lack of information" factor, item EU05 showed high intercorrelations of > .7 with three other items. It was therefore decided to drop this item to avoid redundancy. To identify the most informative items, the remaining items were then submitted to graded response models. Separate models were fitted for each factor. Parameters are given in Table S6.

For the "lack of information" items, *CFI* = 0.978 and *TLI* = 0.969 indicated a good model fit for the graded response model, however *RMSEA* = .093 [.068, .118] indicated a poor fit. Item fit was good for all items with all *p* > .001. Items EU08 and EU11 showed a maximum of item information < 1 and were thus excluded (Coelho et al., 2020). Item EU17 showed low discrimination in comparison to the others and was also excluded. The model with the reduced dataset showed better model fit than before, which was now considered to be acceptable, *CFI* = 0.996, *TLI* = 0.991, *RMSEA* = .061 [0, 0.12]. The remaining items were EU03, EU04, EU07, EU09, and EU10 (α = .901).

Model fit for the graded response model for the "unclear consequences" items was good, *CFI* = 0.987, *TLI* = 0.983, *RMSEA* = .057 [.031, .081]. Item fit was also good for all items with all p > .001. There were five items with a maximum of item information < 1. These were excluded and the model was fitted to the reduced dataset. Model fit was further improved by this, *CFI* > 0.999, *TLI* > 0.999. *RMSEA* = 0 [0, .117]. The remaining items all showed good discrimination and item information. The remaining items were EU01, EU02, EU13, and EU14 (α = .850).

For "normative uncertainty," model fit was acceptable, CFI = 0.981, TLI = 0.976, RMSEA = .064 [.043, .085]. Item fit was good for all items with p > .001. Five items showed a maximum of item information < 1, with NU01, NU11, and NU13 showing the lowest values. It was decided to

keep items NU02 and NU04 because their maximum of item information was > 0.9 and we wanted the "normative uncertainty" factor to have a broad coverage of the underlying construct. We reran the model without NU01, NU11, and NU13. This negatively affected model fit, *CFI* = 0.978, *TLI* = 0.967, *RMSEA* = .085 [.055, .115], but item fit was still given for all items with p > .001. Maximum of item information for NU02 was now close to 1 with 0.993. That of NU04 was at 0.901. The remaining items were NU02, NU04, NU05, NU07, NU08, NU10, and NU14 (α = .860).

Test information and item information curves for the selected items are presented in Figure S1. For these items, we also plotted category characteristics curves to investigate the probabilities of responding to each of the response categories at various levels of the underlying latent variable. This allowed us to estimate whether our choice of a seven-point Likert-scale was appropriate or whether another number of points should be selected. For most items, the category characteristics curves suggested that items might be suited to differentiate between six, but not seven response categories (Figure S2).

Appendix B: Supplementary Tables and Figures

Table S1

Item Pool That Was Submitted to Exploratory Factor Analysis in Study 1

Item	Translation	German original
EU01	It is difficult to judge the action because I do not know what doing or not doing the action would lead to.	Es ist schwierig die Handlung zu beurteilen, weil ich nicht weiß, wozu es führen würde, wenn man die Handlung tut oder unterlässt.
EU02	I would need to know more about the consequences of doing or not doing the action in order to judge its appropriateness.	Ich müsste mehr über die Konsequenzen des Tuns oder Unterlassens der Handlung wissen, um ihre Angemessenheit beurteilen zu können.
EU03	One would need more information to be able to judge the situation clearly.	Man bräuchte mehr Informationen, um die Situation klar beurteilen zu können.
EU04	It is difficult to make a moral decision because I would not know all the relevant facts in the situation.	Es ist schwierig eine moralische Entscheidung zu treffen, weil ich in der Situation nicht alle relevanten Fakten kennen würde.
EU05	I lack important information to be able to judge the action.	Mir fehlen wichtige Informationen um die Handlung beurteilen zu können.
EU06	It would be easier to make a moral decision if I knew more precisely what effects doing or not doing the action would have.	Es wäre leichter, eine moralische Entscheidung zu treffen, wenn ich genauer wüsste, welche Folgen das Tun oder Unterlassen der Handlung hätte.
EU07	I would need more facts to make a decision in this situation.	Ich würde mehr Fakten brauchen um in dieser Situation eine Entscheidung treffen zu können.
EU08	It is difficult to make a moral decision because I don't know what has led to the situation.	Es ist schwierig eine moralische Entscheidung zu treffen, weil ich nicht weiß, was zu der Situation geführt hat.
EU09	Since I do not know everything about the situation, it is difficult to judge the action from a moral point of view.	Da ich nicht alles über die Situation weiß, ist es schwierig, die Handlung aus moralischer Sicht zu beurteilen.
EU10	The moral evaluation of the action is difficult because some facts are unclear.	Die moralische Beurteilung der Handlung ist deshalb so schwierig, weil manche Fakten unklar sind.
EU11	I have all the relevant facts to judge the appropriateness of the action. (reversed)	Ich habe alle relevanten Fakten, um die Angemessenheit der Handlung beurteilen zu können.
EU12	It is foreseeable what morally significant consequences it would have if one does or does not do the action. (reversed)	Es ist vorhersehbar, welche moralisch bedeutsamen Folgen es hätte, wenn man die Handlung tut oder unterlässt.

Item	Translation	German original
EU13	The moral evaluation of the action is so difficult because some consequences are unclear.	Die moralische Bewertung der Handlung ist deshalb so schwierig, weil manche Konsequenzen unklar sind.
EU14	Since I do not know in advance the exact consequences of doing or not doing the action, it is difficult to judge the action from a moral point of view.	Da ich im Vorhinein nicht die genauen Folgen des Tuns oder Unterlassens der Handlung kenne, ist es schwierig, die Handlung aus moralischer Sicht zu beurteilen.
EU15	I am not sure what exactly the consequences of doing or not doing the action would be.	Ich bin nicht sicher, was genau die Folgen des Tuns oder Unterlassens der Handlung wären.
EU16	I am not able to predict the morally significant consequences of doing or not doing the action.	Es ist mir nicht möglich, die moralisch bedeutsamen Folgen des Tuns oder Unterlassens der Handlung vorherzusagen.
EU17	It is difficult to judge the action because the background of the situation is not clear.	Es ist schwierig, die Handlung zu beurteilen, weil die Hintergründe der Situation nicht klar sind.
NU01	Even if I knew everything about the situation, there would be significant doubts as to how to evaluate it morally.	Selbst wenn ich alles über die Situation wüsste, würden große Restzweifel bleiben, wie sie moralisch zu bewerten ist.
NU02	I do not know which of the arguments for or against the appropriateness of the action should be the decisive one.	Ich weiß nicht, welches der Argumente für oder gegen die Angemessenheit der Handlung das entscheidende sein sollte.
NU03	Even if the consequences of doing or not doing the action had been known beforehand, I would be uncertain how to evaluate them morally.	Selbst wenn die Folgen des Tuns oder Unterlassens der Handlung vorher bekannt wären, wäre ich unsicher, wie diese moralisch zu bewerten wären.
NU04	I am unsure how much I should consider the various aspects of this situation in my judgment.	Ich bin unsicher, wie stark ich die verschiedenen Aspekte dieser Situation in meinem Urteil berücksichtigen sollte.
NU05	I have different moral views that lead to different evaluations of this action.	Ich habe verschiedene moralische Überzeugungen, die zu unterschiedlichen Bewertungen dieser Handlung führen.
NU06	I am unsure how to evaluate the consequences of doing or not doing the action from a moral point of view.	Ich bin unsicher, als wie schwerwiegend die Folgen des Tuns oder Unterlassens der Handlung aus moralischer Sicht zu bewerten wären.
NU07	Depending on the perspective from which I think about the situation, I arrive at different evaluations of the action.	Je nachdem, aus welcher Perspektive ich über die Situation nachdenke, komme ich zu unterschiedlichen Bewertungen der Handlung.

ltem	Translation	German original
NOOD	adhere to when judging the action.	Moralvorstellung ich mich bei der Beurteilung der Handlung halten soll.
NU09	With some aspects of the situation I am unsure whether I should consider them in my moral judgment.	Bei manchen Aspekten der Situation bin ich unsicher, ob ich sie in meinem moralischen Urteil berücksichtigen sollte.
NU10	Conflicting arguments come to my mind when judging this action.	Beim Beurteilen dieser Handlung kommen mir widersprüchliche Argumente in den Kopf.
NU11	It is possible that I am wrong in the moral assessment of the situation.	Es ist möglich, dass ich in der moralischen Bewertung der Situation falsch liege.
NU12	Some aspects of the situation make me wonder if they have any bearing on how good or bad the action is morally.	Bei manchen Aspekten der Situation frage ich mich, ob sie einen Einfluss darauf haben, wie gut oder schlecht die Handlung moralisch ist.
NU13	My moral views do not fit together in this situation.	Meine Moralvorstellungen passen in dieser Situation nicht zusammen.
NU14	I am unsure how to apply my moral concepts to this situation.	Ich bin unsicher, wie ich meine Moralvorstellungen auf diese Situation anwenden soll.
NU15	I wonder which aspects of the situation are relevant to their moral evaluation.	Ich frage mich, welche Aspekte der Situation relevant für ihre moralische Bewertung sind.
NU16	I know which moral principles are important in this situation.	Ich weiß, welche moralischen Prinzipien in dieser Situation wichtig sind.
NU17	In this case, I am convinced that my moral principles lead to the correct assessment of the action.	In diesem Fall bin ich überzeugt, dass meine Moralvorstellungen zur richtigen Beurteilung der Handlung führen.
NU18	I am convinced that I am morally correct in my assessment of the situation, even if I do not know all the facts about the situation.	Ich bin überzeugt, dass ich die Situation moralisch richtig einschätze, auch wenn ich nicht alle Fakten über die Situation kenne.
NU19	I know which aspects of the situation are morally relevant.	Ich weiß, welche Aspekte der Situation moralisch bedeutsam sind.
NU20	Even if I had doubts about some of the facts, my moral assessment of the situation is clear.	Selbst wenn ich Zweifel bezüglich mancher Fakten hätte, ist meine moralische Einschätzung der Situation ganz klar.
NU21	I am unsure which aspects of the action are morally significant.	Ich bin unsicher, welche Aspekte der Handlung aus moralischer Sicht bedeutsam sind.

Item	Translation	German original
NU22	I am unsure whether doing or not doing the act is a question of morality.	Ich bin unsicher, ob das Tun oder Unterlassen der Handlung eine Frage der Moral ist.
NU23	It is unclear which aspects of the situation are important for the moral judgment.	Es ist unklar, welche Aspekte der Situation wichtig für die moralische Einschätzung sind.
NU24	I am uncertain whether it makes a moral difference whether one does or refrains from the action.	Ich bin unsicher, ob es moralisch einen Unterschied macht, ob man die Handlung tut oder unterlässt.

Fixed and Random Effects for Predicting Lack of Information Score From Manipulations of

		95% CI		
Fixed effects	b	Lower limit	Upper limit	R^2_{partial}
Intercept	2.275	2.156	2.393	-
Empirical uncertainty manipulation	1.360	1.247	1.474	10.2%
Normative uncertainty manipulation	0.815	0.692	0.937	3.9%
Setting (Oil)	0.723	0.624	0.822	6.1%
Emp. × norm. uncertainty manipulation	-1.335	-1.496	-1.174	5.2%
Random effects	Variance	Corre	lations	
Participant		Intercept	Norm. unc. manipulation	
Random intercept	0.623			
Random slope: Norm. unc. manipulation	0.217	.19		
Randol slope: Setting (Oil)	0.350	15	80	
Residual	1.349			

Empirical and Normative Uncertainty in Study 2

Note. N = 402, 3216 observations, $R^{2}_{fixed} = 14.62\%$, $R^{2}_{total} = 46.36\%$.

Fixed and Random Effects for Predicting Lack of Information Score From Manipulations of

		95%	% CI	
Fixed effects	b	Lower limit	Upper limit	R^2_{partial}
Intercept	5.349	5.172	5.526	-
Empirical uncertainty manipulation	0.361	0.256	0.467	1.0%
Normative uncertainty manipulation	-0.104	-0.212	0.004	0.1%
Setting (Oil)	0.186	0.106	0.266	0.6%
Certainty	-0.487	-0.510	-0.463	32.2%
Emp. × nor. unc. Manipulation	-0.163	-0.307	-0.019	0.1%

Empirical and Normative Uncertainty, Controlling for Certainty, in Study 2

Random effects	Variance	Corre	elations	
Participant		Intercept	Norm. unc. manipulation	
Random intercept	0.374			
Random slope: Norm. unc. manipulation	0.094	.18		
Randol slope: Setting (Oil)	0.138	03	32	
Residual	0.920			

Note. N = 402, 3216 observations, $R^{2}_{fixed} = 41.46\%$, $R^{2}_{total} = 62.05\%$.

Fixed and Random Effects for Predicting Unclear Consequences Score From Manipulations of

		959	% CI	
Fixed effects	b	Lower limit	Upper limit	$R^2_{partial}$
Intercept	2.190	2.072	2.307	-
Empirical uncertainty manipulation	1.277	1.163	1.390	9.2%
Normative uncertainty manipulation	0.867	0.745	0.989	4.4%
Setting (Oil)	0.657	0.560	0.755	5.1%
Emp. × norm. uncertainty manipulation	-1.281	-1.442	-1.121	4.8%
Random effects	Variance	Corre	lations	
Participant		Intercept	Norm. unc. manipulation	
Random intercept	0.601			
Random slope: Norm. unc. manipulation	0.214	.21		
Randol slope: Setting (Oil)	0.314	10	81	
Residual	1.349			

Empirical and Normative Uncertainty in Study 2

Note. N = 402, 3216 observations, $R^{2}_{fixed} = 13.17\%$, $R^{2}_{total} = 45.32\%$.

Fixed and Random Effects for Predicting Unclear Consequences Score From Manipulations of

		95%	6 CI	
Fixed effects	b	Lower limit	Upper limit	$R^2_{partial}$
Intercept	5.236	5.059	5.413	-
Empirical uncertainty manipulation	0.287	0.179	0.396	0.6%
Normative uncertainty manipulation	-0.043	-0.150	0.065	0.0%
Setting (Oil)	0.126	0.052	0.199	0.3%
Certainty	-0.482	-0.506	-0.459	31.6%
Emp. × nor. unc. Manipulation	-0.120	-0.269	0.028	0.1%
Random effects	Variance	Correlation		
Participant		Intercept		
Random intercept	0.361			
Random slope: Norm. unc. manipulation	0.021	.84		
Residual	0.986			

Empirical and Normative Uncertainty, Controlling for Certainty, in Study 2

Note. N = 402, 3216 observations, $R^2_{\text{fixed}} = 39.91\%$, $R^2_{\text{total}} = 58.61\%$.

_	
ø	
S	
Ð	
0	
a	

a
2
ž
e
8
3
te
Ľ
e l
¢
3
Ļ,
j,
₹
2
20
Ø
7
Ш
7
ĕ
i,
5
Ť
A
Ч
717
T Apn.
Study 1
n Study 1
om Study 1
From Study 1
s From Study 1
ites From Study 1
nates From Study 1
timates From Study 1
Estimates From Study 1
g Estimates From Study 1
ing Estimates From Study 1
eling Estimates From Study 1
deling Estimates From Study 1
Aodeling Estimates From Study 1
Modeling Estimates From Study 1
se Modeling Estimates From Study 1
onse Modeling Estimates From Study 1
ponse Modeling Estimates From Study 1
esponse Modeling Estimates From Study 1
Response Modeling Estimates From Study 1
ed Response Modeling Estimates From Study 1
ded Response Modeling Estimates From Study 1

Table S6															
Graded R	esponse N	Aodeling I	Estimates	From Stuc	iy 1 After	Final EFA	and After	Further It	em Remo	val					
				Befor	e item rer	noval					After	item rem	oval		
Factor	ltem	a	b_1	b_2	b_3	b_4	b_5	b_6	a	b_1	b_2	b_3	b_4	b_5	b_6
Info	EU03	3.06	-0.98	-0.25	0.11	0.37	1.00	1.53	3.15	-0.96	-0.25	0.10	0.36	0.99	1.51
	EU04	2.97	-0.95	-0.34	0.19	0.39	1.13	1.85	2.88	-0.96	-0.36	0.17	0.38	1.14	1.87
	EU07	2.82	-0.93	-0.15	0.33	0.63	1.22	1.66	2.70	-0.93	-0.14	0.33	0.64	1.23	1.68
	EU08	1.01	-0.51	0.59	0.97	1.42	2.11	3.11	ı	ı	ı	ı	·	ı	ı
	EU09	3.22	-0.99	-0.24	0.24	0.68	1.39	2.06	3.31	-0.97	-0.23	0.23	0.67	1.37	2.04
	EU10	2.52	-1.04	-0.25	0.13	0.48	0.96	1.87	2.72	-1.01	-0.26	0.12	0.47	0.93	1.81
	EU11	1.49	-1.85	-0.76	-0.03	0.56	1.30	2.14	ı	ı	ı	ı	ı	ı	ı
	EU17	1.80	-0.84	0.05	0.58	0.86	1.51	2.36	ı	ı	ı	ı	ı	ı	ı
Cons	EU01	2.72	-1.32	-0.66	-0.14	0.12	0.65	1.43	3.13	-1.26	-0.64	-0.14	0.10	0.61	1.39
	EU02	2.24	-1.50	-0.57	-0.19	0.13	0.87	1.60	2.11	-1.54	-0.58	-0.19	0.13	0.89	1.65
	EU06	1.73	-1.79	-1.24	-0.79	-0.46	0.32	1.10	ı	ı	ı	ı	ı	ı	ı
	EU12	0.86	-2.58	-1.48	0.17	0.83	1.80	3.07	ı	ı	ı	ı	ı	ı	ı
	EU13	2.24	-1.36	-0.72	-0.40	-0.10	0.39	1.14	2.35	-1.34	-0.71	-0.40	-0.11	0.38	1.12
	EU14	2.24	-1.36	-0.52	-0.01	0.37	1.07	1.89	2.12	-1.39	-0.53	-0.01	0.38	1.10	1.94
	EU15	1.72	-1.70	-0.66	-0.28	0.13	0.86	1.98	·	ı	·	·		·	ı
	EU16	1.67	-1.72	-0.65	0.01	0.51	1.15	1.87	·	ı	·	·		·	ı
	90NN	1.21	-2.49	-1.24	-0.38	0.06	1.00	2.52	ı	ı	ı	ı	ı	ı	ı

				Befor	e item rer	noval					After	item rem	oval		
Factor	ltem	a	b_1	b_2	b_3	b_4	b_5	b_6	a	b_1	b_2	b_3	b_4	b_5	b_6
Norm	NU01	1.46	-1.18	-0.41	0.09	0.52	1.42	2.40	ı	1	ı	ı	ı	I	
	NU02	1.68	-1.42	-0.61	-0.20	0.23	0.91	1.84	1.75	-1.39	-0.61	-0.21	0.21	0.88	1.81
	NU04	1.70	-2.11	-1.10	-0.66	-0.18	0.65	1.89	1.67	-2.11	-1.11	-0.66	-0.18	0.65	1.92
	NU05	2.02	-1.83	-1.09	-0.54	-0.25	0.48	1.23	1.99	-1.83	-1.11	-0.56	-0.27	0.47	1.24
	NU07	1.85	-2.41	-1.50	-0.96	-0.52	0.12	0.92	1.93	-2.35	-1.47	-0.95	-0.52	0.11	0.89
	NU08	2.37	-1.41	-0.66	-0.24	0.16	0.83	1.64	2.30	-1.42	-0.68	-0.26	0.15	0.84	1.66
	NU10	2.14	-2.16	-1.48	-0.89	-0.65	0.06	0.69	2.26	-2.09	-1.43	-0.87	-0.64	0.05	0.67
	NU11	1.12	-2.80	-1.72	-0.99	-0.27	0.82	1.95	I	I	·	·	ı	ı	I
	NU13	1.31	-2.03	-0.89	-0.21	0.31	1.20	2.14	I	ı	ı	ı	,	ı	ı
	NU14	1.92	-1.42	-0.74	-0.22	0.18	1.15	2.01	1.86	-1.43	-0.75	-0.23	0.18	1.17	2.05
			,												

Note. Models were fitted to items of each factor separately. a: discrimination, b_j: threshold, Info: "lack of information", Cons: "unclear consequences",

Norm: "normative uncertainty".

Test Information (A) and Item Information (B) for Each Extracted Factor After Item Removal


Figure S2

Category Characteristics Curves for the Final Item Selection After Study 1



Note. Info: "lack of information"; Cons: "unclear consequences"; Norm: "normative uncertainty".

A-2 Manuscript 2 – Metaethical Intuitions in Lay Concepts of Normative Uncertainty

Copyright © 2023 by Taylor & Francis. Reproduced with permission.

Theisen, M. (2023). Metaethical intuitions in lay concepts of normative uncertainty. *Philosophical Psychology*. Advance online publication. https://doi.org/10.1080/09515089.2023.2174422

Metaethical Intuitions in Lay Concepts of Normative Uncertainty

Maximilian Theisen¹

¹Department of Psychology, Heidelberg University, Heidelberg, Germany

Author Note

The author would like to thank Julia M. Lemp, James Andow, and the anonymous referees for their insightful comments on different versions of this article. This work was supported by an individual fellowship of the Baden-Württemberg *Landesgraduiertenförderung* Program from Heidelberg University, Germany. The data that support the findings of this paper are openly available on OSF at https://osf.io/pxrgh. The author reports there are no competing interests to declare.

Correspondence concerning this article should be addressed to Maximilian Theisen, Department of Psychology, Heidelberg University, Hauptstraße 47-51, 69120 Heidelberg, Germany. Email: maximilian.theisen@psychologie.uni-heidelberg.de

Abstract

Even if we know all relevant descriptive facts about an act, we can still be uncertain about its moral acceptability. Most literature on how to act under such normative uncertainty operates on moral realism, the metaethical view that there are objective moral facts. Lay people largely report anti-realist intuitions, which poses the question of how these intuitions affect their interpretation and handling of normative uncertainty. Results from two quasi-experimental studies (total *N* = 365) revealed that most people did not interpret normative uncertainty as referring to objective moral facts but rather as uncertainty regarding one's own view, uncertainty regarding the culturally accepted view or as the result of ambivalence. Especially the anti-realist majority of participants interpreted normative uncertainty different to how it is described in the literature on choice under normative uncertainty. Metaethical views were also associated with lay peoples' choice of uncertainty reduction strategies and with assumptions about the intended aim of such strategies. The current findings suggest that empirical investigations of normative uncertainty might benefit from considering folk metaethical pluralism, as the lay public largely disagrees with the metaethical assumptions underlying the current discourse on choice under normative uncertainty.

Keywords: moral uncertainty; normative uncertainty; ambivalence; metaethics; moral psychology; non-cognitivism

Introduction

When we are uncertain about the morality of an act, this uncertainty can be grounded in ignorance of the relevant descriptive facts, such as the tangible consequences the act would bring. It can also be grounded in uncertainty regarding the moral evaluation of these facts, such as whether a given consequence would be morally good or bad. This latter form of uncertainty has been called normative uncertainty (Bykvist, 2017).

In the last two decades, there has been a growing body of literature addressing both metaethical (e.g., Bykvist & Olson, 2017; Eriksson & Francén Olinder, 2016; Sepielli, 2012; Smith, 2002; Staffel, 2019) and normative aspects of normative uncertainty (e.g., Bykvist, 2017; Lockhart, 2000; MacAskill et al., 2020; Moller, 2011; Sepielli, 2017). From a normative perspective, philosophers have investigated what the most rational way to act under normative uncertainty would be.¹ This approach has been called uncertaintism (Harman, 2015). Some uncertaintists have argued that we should act according to the view that we give most credence to (Gustafsson & Torpman, 2014; Gustafsson, 2022). Others have argued for moral hedging – a view according to which we should take into account all moral views that we give at least some credence to and consider both our credences and the moral value each view would ascribe to the act in question (e.g., Bykvist, 2017; Lockhart, 2000).

From a psychological perspective, it needs to be asked how non-philosophers act under normative uncertainty.² How do they weight conflicting moral views, arguments, and ideas when

¹ This focus on rationality has recently been criticized by Rosenthal (2021) who calls for a moral rather than a rational evaluation of choice under normative uncertainty.

² First steps towards an empirical investigation of normative uncertainty have recently been taken by Costa-Gomes and Schönegger (2022), and Jabarian (2020). For recent work on uncertainty in moral judgments without explicit reference to uncertaintism, see for example Alsaad (2021), Alsaad et al. (2021), Guzmán et al. (2022), and Mata (2019). they have to make decisions? Investigating these issues requires that researchers have some understanding of the terms in which lay people think about these concepts. However, some expert notions of normative uncertainty operate on metaethical ideas that lay people by and large reject. This article aims to reveal the assumptions that are inherent to large parts of the uncertaintism discourse, how these assumptions can conflict with metaethical intuitions in lay people, and how this is relevant for empirical investigations of normative uncertainty.

Metaethical Assumptions in the Uncertaintism Discourse

Metaethics is concerned with the nature of morality itself and the characteristics of moral thought and language. In this section, I will show how different metaethical views affect what is meant by the term *normative uncertainty* and which of these views have prevailed in the uncertaintism discourse so far. I will limit my discussion to views that are relevant both to analyses of normative uncertainty and to folk metaethics. This will entail the distinction of cognitivism and non-cognitivism, and a discussion of realism and different forms of relativism.³ Table 1 gives an overview of the discussed concepts.

Cognitivism and Non-Cognitivism

One rather explicitly made assumption in the uncertaintism literature is that of moral cognitivism (Makins, 2021; Nissan-Rozen, 2015; Sepielli, 2012). Cognitivism entails that moral judgments are beliefs and that moral sentences are expressions of those beliefs that can be true or false (truth-apt; van Roojen, 2018). Non-cognitivism denies both aspects of this

³ Among others, this will exclude error theorist accounts of morality (Joyce, 2011;

Mackie, 1977). Error theorists usually assume that moral language intends to state moral facts, but that there are no such facts, which would make all moral propositions untrue. This metaethical view leaves no room for uncertainty about the truth of a given moral proposition. Since error theory is also endorsed by only a small fraction of lay people (Pölzler & Wright, 2020; Wright & Pölzler, 2021), this view was not featured in the current studies. characterization. On a psychological level, non-cognitivism states that moral judgments are not beliefs but non-cognitive states, such as desires, preferences, emotions, or intentions (Eriksson & Francén Olinder, 2016). On a semantic level, non-cognitivism states that moral sentences are not truth-apt – they cannot be true or false. Usually, non-cognitivists also deny the existence of moral facts altogether (Makins, 2021; van Roojen, 2018).

While cognitivists understand normative uncertainty to refer to subjective or objective moral facts and describe it as the strength of a belief concerning the morality of an act, the compatibility of non-cognitivism and normative uncertainty is still an open problem. According to Smith (2002), non-cognitive states do not possess the structural features to account simultaneously for the strength of a judgment and its certitude. As a response, some have proposed modified accounts of non-cognitivism, either by assuming the involvement of auxiliary beliefs (Eriksson & Francén Olinder, 2016; Lenman, 2003), or by introducing higher order non-cognitive states (Beddor, 2020; Sepielli, 2012; Staffel, 2019). However, several concerns regarding these accounts have been raised (Björkholm et al., 2021; Bykvist & Olson, 2017; Lam, 2020; Makins, 2021).⁴

⁴ Some of those who try to accommodate normative uncertainty within non-cognitivism do so from a quasi-realist stance (e.g., Beddor, 2020; Lam, 2020). Quasi-realism is a variant of non-cognitivism that aims to make sense of ordinary moral practice, such as speaking of moral facts, from a non-cognitivist perspective (Blackburn, 1993; Gibbard, 1990). The complexity of quasi-realism makes it an unlikely candidate for describing folk metaethical intuitions. Further, boundaries between traditional non-cognitivist and quasi-realist analyses of normative uncertainty are often blurred (see for example Beddor, 2020 and Sepielli, 2012). Therefore, quasi-realism will not be discussed as a separate metaethical view in this paper. Recent quasirealist discussions of normative uncertainty or the associated concept of moral fallibility can be found in Beddor (2020), Björkholm et al. (2021), Egan (2007), Köhler (2015) or Lam (2020).

Makins (2021) proposes a different approach. He argues that pure non-cognitive states do possess the structural complexity that is observed in moral judgments, since non-cognitive states vary not only in strength and temporal stability but also in internal consistency. Drawing from psychological literature on attitudes, Makins continues by showing that what has been described as normative uncertainty can also be interpreted as the result of ambivalence. In contrast to other proposals, this interpretation of normative uncertainty does not depend on the assumption of auxiliary beliefs or additional non-cognitive states. Illustrated for the seminal Trolley dilemma (Foot, 1967), a cognitivist might explain their normative uncertainty as a divided credence in the moral view that, in this situation, it would be right to minimize the number of casualties and the view that it would be right to refrain from active killing. A non-cognitivist might interpret uncertainty as ambivalence resulting from conflicting desires (or intentions, etc.) to minimize the number of casualties and to refrain from active killing.⁵

The debate concerning the compatibility of normative uncertainty and non-cognitivism shows that underlying metaethical assumptions are relevant to how normative uncertainty is conceptualized. From a cognitivist view, normative uncertainty is uncertainty about whether a moral proposition is true or false and can be described as the strength of a belief. From a non-

⁵ Recently, Robinson and Steele (2022) have proposed to understand reasoning under normative uncertainty from a non-cognitivist perspective as reasoning involving multiple objectives. Similar to Makins, they understand normative uncertainty to be the result of conflicting non-cognitive states. Unlike Makins, the multi-objective problem approach, in its simplest form, does not differentiate between being torn between two moral perspectives and being certain in a compromise between those perspectives. Robinson and Steele call this the differentiation problem. However, Robinson and Steele's approach offers a greater formalization of decision-making under normative uncertainty than Makins', including a representation of different meta-normative theories, such as moral hedging. cognitivist view, normative uncertainty may describe ambivalence resulting from conflicting evaluations of an act, without assuming the existence of moral facts.

Realism, Cultural Relativism, and Individual Relativism

When assuming that normative uncertainty refers to moral facts, it needs to be asked what constitutes them. Within cognitivism, there are various views on this matter. One of these is moral realism. Moral realism is the notion that moral properties, such as being morally right or wrong, are objective features of the world.⁶ This entails that there are moral facts that go beyond how individuals or cultures think about the respective issues (Brink, 1989; Colebrook, 2021; Railton, 1986; Shafer-Landau, 2003). Some state that such objective moral features are reducible to non-moral features (naturalism; Lutz & Lenman, 2021), others deny this (nonnaturalism; Moore, 1903). This distinction can affect how people are understood to learn about objective moral facts (Ridge, 2019). Normative uncertainty from a realist perspective refers to whether one's moral beliefs accurately reflect these objective facts.

Another view on what constitutes moral facts is relativism. Relativism denies that there are any objective moral facts and instead assumes that such facts completely depend on how cultures (cultural relativism; Tilley, 2000) or individuals think about them (individual relativism; Ewing, 1948). For cultural relativists, normative uncertainty would be uncertainty about whether one's moral beliefs are congruent with the dominant moral views or values in one's culture or society. For individual relativists, normative uncertainty could be understood as uncertainty about what one's moral beliefs actually are (see also Petrocelli et al., 2007) or whether they accurately reflect one's deeply held moral values (Moller, 2011).

⁶ For the present purpose, it is sufficient to treat realism and objectivism as equivalent. Note, however, that there are fine-grained differences between both notions; for example, see Pendlebury (2011). In the uncertaintism literature, the existence of moral facts is usually assumed (MacAskill et al., 2020; Makins, 2021; Nissan-Rozen, 2015; Rowland, 2021) but uncertaintists seldomly state which account of moral facts they refer to. However, there are some exceptions to this. MacAskill et al. (2020) asserted that their work on uncertaintism is based on "some form of moral realism" (p. 147) and that they deem both non-cognitivism and error theory as incompatible with normative uncertainty. Similarly, Rowland (2021) describes the uncertaintism discourse as referring to what is "objectively" morally required (p. 142) and Makins (2021) interprets the terms that are used in uncertaintism frameworks to "betray an underlying presupposition in favour of metaethical realism" (p. 2).

Tarsney (2021) points out that such realist assumptions may have even motivated the uncertaintism project. In his review of Weatherson (2019), Tarsney attributes Weatherson's critique of the uncertaintism discourse in part to Weatherson's alleged anti-realism. Tarsney states that it is easier to make sense of uncertaintism when one assumes a "metaphysical symmetry between normative and empirical facts", and he argues further that "the felt need for a theory of decision-making under normative uncertainty" might be more pronounced from a realist than an anti-realist perspective (p. 1021).

Sepielli (2017) makes a similar argument. Uncertaintists usually discuss normative uncertainty in terms of credences and subjective probabilities (Lockhart, 2000; MacAskill et al., 2020). This is what Tarnsey might see as a reference to a metaphysical symmetry between the empirical and the normative. Sepielli (2017) specifies that these subjective probabilities should, in his view, be understood as epistemic probabilities of "objective moral claims" (p. 101).

Even though the idea of uncertaintism might be rooted in a realist perspective, Moller (2011) shows how arguments about what one ought to do under normative uncertainty do not necessarily require that moral facts are objective. He proposes an account of moral error, and thus the possibility of uncertainty regarding which moral view is right, where moral error describes "inconsistencies between our practices and beliefs, and our own deepest values" (p. 434). From this individual relativist perspective, uncertaintism might be motivated by the desire not to act against one's own moral values.

In this section, I showed that different metaethical perspectives suggest different interpretations of the term normative uncertainty. This is most obvious for the distinction of cognitivism and non-cognitivism, where normative uncertainty is either interpreted in terms of beliefs or in non-cognitive terms, such as ambivalence. Exploring further distinctions within cognitivism has shown that moral beliefs can either refer to objective or subjective facts. I further hope to have shown that the discourse on how to act under normative uncertainty has largely been characterized, and may even have been motivated, by realist notions of morality. In the next section, I will address how lay people think about metaethical matters to see whether their views conflict with uncertaintists' realist assumptions.

Metaethical Intuitions in Lay People

In the past two decades, non-philosophers' intuitions regarding the nature of morality and moral judgments have become the object of empirical investigation (for reviews, see Colebrook, 2021; Pölzler & Wright, 2019). A variety of paradigms have evolved to measure these intuitions, with researchers largely relying on semi-implicit methods, such as disagreement or truth-aptness tasks (e.g., Beebe et al., 2015; Beebe & Sackris, 2016; Goodwin & Darley, 2008, 2012; Sarkissian et al., 2011; Wright & Pölzler, 2021), but purely implicit (Wagner et al., 2021; Zijlstra, 2021) and explicit methods have also been proposed (Pölzler & Wright, 2020; Zijlstra, 2019). Results from these studies paint the picture of a dominantly anti-realist folk (Colebrook, 2021; Pölzler & Wright, 2019, 2020). However, lay people's metaethical intuitions have been shown to vary between moral issues, with some issues having a higher chance of eliciting certain metaethical judgments than others (metaethical pluralism). Additionally, substantial variance in reported metaethical intuitions between persons and between measurement techniques has been documented (Colebrook, 2021; Pölzler & Wright, 2020).

Study 1

The evidence so far suggests that lay people are far from being consistent moral realists. However, large parts of the uncertaintism discourse seem to assume realism. In this study, I investigated how lay people interpret normative uncertainty depending on their own metaethical intuitions. I hypothesized that only realists would interpret uncertainty in a way that is consistent with the uncertaintism discourse.⁷ For cultural relativism, I expected participants to interpret uncertainty as uncertainty regarding the moral view that is held in their culture. For individual relativism, I expected participants to interpret uncertainty as uncertainty regarding one's own moral values. Last, for non-cognitivist participants, I expected that uncertainty would be interpreted as ambivalence.

Besides interpretations of normative uncertainty, I also hypothesized that strategies to resolve such uncertainty would depend on participants' metaethical presumptions. Whilst realist participants should be interested in learning new objective moral information (MacAskill et al., 2020), cultural relativists should be interested in learning how people in their culture think about the respective issue. In contrast, individual relativists should be concerned with discovering what their true moral values in the situation are.

For non-cognitivists, I assumed that they would indicate that the uncertainty cannot be deliberately resolved, since their uncertainty is not defined as ignorance of a fact and can therefore not be resolved by learning that respective fact. Even though various lay strategies for resolving ambivalence have been described in the literature, it was questionable whether these would be suitable as response options for the current studies. This was either because these

⁷ For pragmatic reasons, I will use the terms "realists", "non-cognitivists" etc. to refer to the metaethical judgments of the participants. However, given the high level of intraindividual variance in metaethical judgments, these terms should be understood as depending on the respective issue rather than as domain-general personality traits.

strategies address the aversiveness of ambivalence rather than the ambivalence itself (e.g., emotion-focused strategies), they are not deliberate (e.g., biased information processing), or they rely on acquiring new information which would resemble the cognitivist strategies that were already considered in the study material (van Harreveld et al., 2009). Additionally, there were positive reasons to expect that non-cognitivists would indicate that their uncertainty cannot be deliberately resolved, since this is consistent with an understanding of ambivalence as the result of mutually incompatible non-cognitive states (Makins, 2021; Shore, 1990).

Method

Participants

Data was collected online using the *SoSci Survey* software (Leiner, 2022). Participants (N = 207) were recruited via mailing lists from German universities. Six participants were excluded from the analyses, since their data suggested that they had speeded through the assessment, indicated by a relative speed index > 2 (Leiner, 2019). Further 16 participants had to be excluded due to a technical error during data collection and 3 participants failed more than one attention check, leaving N = 182 participants (133 female, 44 male, 3 explicitly unspecified, 1 diverse, and 1 not indicating). The sample size was comparable to that of other studies using similar measurements (Pölzler & Wright, 2020; Wright & Pölzler, 2021). Participants were aged between 18 and 71 (M = 25.08, SD = 9.05). Most of them had a high educational degree, equivalent to A-level (N = 130). N = 43 had a bachelor's degree or higher. Six further participants had completed an apprenticeship. One participant had a medium educational level and one was still in school. As reimbursement, participants either received course credit or could take part in a lottery for two 10€ vouchers.

Design and Material

The study followed a quasi-experimental design that grouped participants according to their pre-existing metaethical views. Participants were presented with six tasks in total. The metaphor task, the disagreement task, the theory task, and the category task built on tasks described in Wagner et al. (2021) and were translated to German by the author. Two further tasks – an uncertainty task and a resolving task – were newly developed for this study.

Each task contained questions about issues of science, morality, social convention, and personal taste (Table 2). The issues were adapted from Wagner et al. (2021), however some were replaced to better fit the culture of the study sample. Each issue was presented in each task except for the issue of rape, which was not presented in the resolving task since it was considered to be too extreme to ask participants to imagine that they would be uncertain regarding the permissibility of this act. In the metaphor, the theory, and the resolving task, one additional trial was presented as attention check. Except for these attention checks, the order of trials was randomized for each task. The order of tasks was as described in the following. A translation of the full task material can be found in the supplements of this article.

Metaethical Tasks. The assessment started with the metaphor task. For each issue, participants were asked to indicate which metaphor would best describe whether there is a correct answer to the respective question and, if so, what it is that makes this answer correct. The metaphor task differentiated between realism (the answer is "out there in the world"), relativism (the answer is "invented or created by individuals or societies"), and non-cognitivism (there is not a correct answer, but any response to the question would just be like someone saying "boo!" or "hurray!").

This task was followed by the disagreement task where participants interpreted two forms of disagreement regarding the presented issues. First, they interpreted disagreement between two members of the same culture, then between two members of different cultures. Participants were asked to indicate whether one of both persons is correct and the other incorrect, whether both are correct, or whether neither of them is correct nor incorrect. Given the response patterns across both forms of disagreement, the disagreement task was able to differentiate between realism, cultural relativism, individual relativism, non-cognitivism, and non-standard views (see Wagner et al., 2021). The last of the metaethical tasks was the theory task. This task was similar to the metaphor task except that the different metaethical views were described without reference to metaphors. The theory task differentiated between realism, relativism, and non-cognitivism.

Uncertainty Task. The goal of the uncertainty task was to reveal how participants interpret normative uncertainty. For each issue, participants were asked to imagine someone that was uncertain about this issue. Asking participants to interpret another person's uncertainty made it possible to involve moral cases that the participants themselves were not necessarily uncertain about. Then, they were asked to indicate which interpretation of such uncertainty they deemed to be the most appropriate/most obvious (German original: "am naheliegendsten"). I hypothesized that the first option would refer to a realist, the second option to cultural relativist, the third option to individual relativist, and the last option to noncognitivist notions of uncertainty.

Imagine that someone was uncertain whether **it is morally acceptable for a state to limit immigration**. We will present you with several interpretations of what it may mean to be uncertain about this question.

Which interpretation seems most appropriate/most obvious to you?

- (1) The person does not know the objectively correct answer to this question.
- (2) The person does not know the answer that is considered correct by society.
- (3) The person does not know their own stance or opinion in this question.
- (4) The person has conflicting feelings, intentions, emotions or attitudes towards this question.

Resolving Task. The resolving task was designed to measure which strategies participants would choose to resolve normative uncertainty. Therefore, it was necessary that they imagined themselves to be uncertain about the described issues. I hypothesized that the response pattern would be the same as in the uncertainty task. The task read as follows: Imagine that you yourself were uncertain whether it is morally acceptable for a state to limit immigration.

How could you best reduce this uncertainty?

- (1) I could look at what science or philosophy have to say about this question.
- (2) I could look at which answer to this question is given in my culture.
- (3) I could try to become clearer about my own opinion on this question.
- (4) This uncertainty could not be deliberately reduced.

Category Task. Lastly, participants were asked to categorize each issue as either a matter of science, morality, social convention, or personal taste. This approach has been proposed in the literature since previous research had suggested that there is dissent in participants about which issues can be considered of moral relevance (e.g., Pölzler & Wright, 2020; Skitka et al., 2021; Wright et al., 2008).

Analysis

To test for associations between two categorical variables (e.g., responses to the theory task and responses to the uncertainty task), Chi squared tests were used. Such tests are based on contingency tables and can inform us whether the distribution of responses to one task differs between responses to another task (e.g., whether distributions of responses to the uncertainty task differ between participants with different metaethical views). In the case of a significant Chi squared test, post hoc comparisons were used to clarify which response categories contribute to the overall effect (Sharpe, 2015). For example, post hoc comparisons were used to test whether realists showed significantly higher rates of realist interpretations of normative uncertainty than the overall sample.

If the expected value of a cell in the contingency table was smaller than five (e.g., less than five realists who interpret normative uncertainty as ambivalence), *p*-values were estimated using a Monte Carlo procedure with 2000 simulations (Hope, 1968). In these cases, degrees of freedom for the Chi squared test are not reported. As an effect size, bias-corrected Cramér's *V* is indicated (Bergsma, 2013). Since the interpretation of Cramér's V depends on its degrees of freedom, these are indicated in parentheses (Cohen, 1988).

For post hoc comparisons, adjusted standardized residuals were calculated which can be compared across contingency tables of varying dimensions and sample sizes (Agresti, 2007; Sharpe, 2015). Due to the high number of comparisons, *p*-values for these post hoc tests were Bonferroni-corrected, which is a conservative correction (Jafari & Ansari-Pour, 2019).

Results

Category Task

Each issue was dominantly categorized as intended (Table 2). However, there was some disagreement regarding the categorization of the issue of rape. Whilst this question had been rated as a moral issue by 94% of participants in Wagner et al. (2021), it was only 76.9% in this study. Similar disagreement was only present in the pyjamas and in the sandals question, which were partly rated as matters of social convention and partly as matters of personal taste. The issues that were newly developed for this study, namely sun and immigration, proved to fit the intended categories well. In the following, the term categories will be used to refer to the classifications made by the participants.

Metaethical Tasks

The distributions of responses to the metaethical tasks across categories and issues can be seen in Figure 1. Whether participants categorized the presented issue as an issue of science, morality, social convention, or personal taste had large effects on how they responded to the metaphor task, $X^2(6) = 681.28$, p < .001, V(2) = .454, the disagreement task, $X^2(12) = 617.37$, p <.001, V(3) = .351, and the theory task, $X^2(6) = 1109.38$, p < .001, $V(2) = .581.^8$ To evaluate the internal consistency of the metaethical tasks, I investigated associations between the metaphor

⁸ If not given as figures, post hoc tables for each chi-square test can be found in the supplements.

task, the disagreement task, and the theory task for moral issues only. The highest association was found between the disagreement and the theory task, $X^2 = 175.73$, p < .001, V(2) = .423 (large effect size), followed by the association between the theory and the metaphor task, $X^2(4) = 94.91$, p < .001, V(2) = .311 (medium effect size), and the association between the metaphor and the disagreement task, $X^2 = 86.57$, p < .001, V(2) = .289 (medium effect size). Uncertainty Task

The distributions of responses to the uncertainty and the resolving task across categories and issues can be seen in Figure 2. Interpretations of uncertainty varied between categories with a large effect size, $X^2(9) = 1222.94$, p < .001, V(3) = .497. For issues categorized as moral, uncertainty was mostly interpreted as ambivalence (34.8%), followed by interpretations of ignorance of one's own view (22.9%), ignorance of the culturally accepted view (22.1%), and ignorance of an objective fact (20.2%).

These interpretations of normative uncertainty depended on the participants' metaethical views. Responses to the uncertainty task were associated with responses to the metaphor task, $X^2(6) = 55.52$, p < .001, V(2) = .230 (medium effect size), the disagreement task, $X^2 = 95.68$, p < .001, V(3) = .244 (medium effect size), and the theory task, $X^2(6) = 110.47$, p < .001, V(2) = .334 (medium effect size). As can be seen in Figure 3, these effects were mainly driven by realist participants showing elevated rates of interpreting normative uncertainty as ignorance of an objective fact. Relativists as identified by the theory task showed elevated rates of interpreting normative uncertainty as ignorance of the culturally accepted view. Last, non-cognitivist responses to the disagreement and the theory task were associated with elevated rates of interpreting normative uncertainty as ambivalence.

Resolving Task

How participants would try to resolve their uncertainty varied between categories with a large effect size, $X^2(9) = 1511.33$, p < .001, V(3) = .587. For moral issues, most participants wanted to increase clarity regarding their own view (47.7%), followed by consulting science or

philosophy (35.1%), looking for an answer in one's culture (13.6%), and indicating that the uncertainty cannot be deliberately resolved (3.6%).

These responses depended on the participants' metaethical views. For moral issues, responses to the resolving task were associated with responses to the metaphor task, $X^2 = 16.06$, p = .017, V(2) = .124 (small effect size), responses to the disagreement task, $X^2 = 37.18$, p < .001, V(3) = .160 (small effect size), and responses to the theory task, $X^2 = 49.04$, p < .001, V(2) = .256 (medium effect size). Across all metaethical tasks, realist participants showed elevated rates of consulting science or philosophy (Figure 4). Associations with the theory task revealed that the rate of participants looking for an answer in their culture was elevated for relativist participants and that non-cognitivists showed elevated rates of wanting to increase clarity regarding their own views.

Responses to the resolving task were also associated with interpretations of normative uncertainty, $X^2 = 45.7$, p < .001, V(3) = .193 (medium effect size). Interpreting normative uncertainty as ignorance of an objective fact was linked to choosing science or philosophy, interpreting it as ignorance of the culturally accepted view was linked to looking for an answer in one's culture, and interpreting it as ambivalence was linked to indicating that the uncertainty cannot be deliberately resolved (Figure 5).

Discussion

The results of Study 1 largely supported the main arguments of this paper. First, only a minority of participants shared a realist notion of normative uncertainty as it has been construed in the uncertaintism discourse. Instead, most participants interpreted normative uncertainty in anti-realist ways. Second, it could be shown that the mismatch between lay and uncertaintist notions of normative uncertainty can be linked to the underlying metaethical assumptions.

Uncertainty Task

The identified associations between the uncertainty and the metaethical tasks were mostly driven by realist and non-cognitivist respondents. The link between non-cognitivism and interpreting normative uncertainty as ambivalence was notable. From a lay perspective, noncognitivist uncertainty might as well have been understood as uncertainty regarding one's own (non-cognitive) mental states, which would have shown up as individual relativist responses to the uncertainty task. Such an interpretation of non-cognitivist uncertainty would conflict with Smith (2002) who argues that non-cognitive states do not possess the structure to simultaneously account for certainty and strength. However, lay peoples' reports of metaethical beliefs are not necessarily internally consistent (Colebrook, 2021), so it was not ruled out that they would respond this way. Yet, the metaethically sound association between non-cognitivism and ambivalence was a rather consistent finding in the current study. This also gives some support to Makins (2021), who proposes to describe non-cognitivist uncertainty as ambivalence. It has to be acknowledged that concurring accounts of non-cognitivist uncertainty, such as those by Eriksson and Francén Olinder (2016), Sepielli (2012), or Lenman (2003), were not represented in this study. However, given the complexity of these accounts, it seems questionable whether they are probable candidates for describing lay concepts of normative uncertainty.

In regard to relativism, post hoc comparisons were largely non-significant. This might be explained by the fact that relativism, as operationalized in Study 1, entails two distinct metaethical notions, namely cultural relativism and individual relativism, but this distinction was only possible in the disagreement task. However, differentiating between both forms of relativism is especially relevant for the current investigation since they imply different interpretations of uncertainty, with one referring to external and the other to internal moral facts. Therefore, despite the inconsistent findings in the current study, there were strong theoretical reasons to further investigate this differentiation. To do so, metaphor and theory task were adapted for Study 2 so that they would distinguish between cultural and individual relativism.

Another worry in regard to the uncertainty task was the high rate of ambivalence interpretations across all metaethical views and across relativists in particular. One explanation for this might be that different forms of anti-realism are psychologically not as distinct as could be expected from a metaethical perspective (Zijlstra, 2019). However, this would not explain why even among realists more than 20% chose this response option. A second explanation might be that the ambivalence response option was stated in a way that would make it not only attractive to non-cognitivist but also to cognitivist participants. Cognitivism does not preclude that uncertainty can be accompanied by conflicting non-cognitive states (e.g., Turp, 2018). If the ambivalence response was understood not to include the denial of truth-aptness then this response option might have been a way for participants to interpret normative uncertainty without taking position on complex metaethical matters. To address these concerns in Study 2, the non-cognitivist response option was revised to explicitly exclude cognitivist readings of normative uncertainty.

One could also argue that some realists might have been drawn towards interpreting normative uncertainty as ignorance of the culturally accepted view or as ignorance regarding their own view. Realism does not preclude to be uncertain about these aspects and some realists might even see them as relevant indicators of objective moral truth (Ayars & Nichols, 2020; MacAskill et al., 2020; Ridge, 2019). Even though such participants would still be uncertain which moral view is objectively correct, they might not have indicated this in their responses to the uncertainty task, which would have weakened the associations between metaethical views and uncertainty interpretations. To address this potential issue, I also revised the cognitivist response options for the uncertainty task to put more emphasis on the aspect of what the uncertainty *ultimately* refers to.

Resolving Task

In regard to resolving normative uncertainty, most participants either chose to increase clarity regarding their own moral views or they chose to consult science or philosophy. These preferences depended on the participants' metaethical views, which was mainly driven by realist respondents showing higher rates of consulting science or philosophy than anti-realists. Overall, associations with the resolving task were smaller than those with the uncertainty task.

There were, of course, also some limitations to the resolving task. It was expected a priori that moral realists would prefer to consult science or philosophy and that individual relativists would prefer to engage in introspective processes to find out what their true moral values are. Yet some realists might ascribe evidential value to mental processes, such as intuition or reasoning (MacAskill et al., 2020; Ridge, 2019), and might therefore try to resolve normative uncertainty in this way rather than consulting the literature. Similarly, individual relativists might try to identify their own moral values by confronting themselves with a variety of ethical alternatives, for example in philosophical discourse or by learning about the experiences of potential victims. As Eriksson and Francén Olinder (2016) put it: "[W]hen trying to rid ourselves of the uncertainty [about our own moral values], we don't direct our attention to our own minds – rather, we direct it to the acts in question and their characteristics" (p. 728).

This suggests that there are two separate aspects to resolving uncertainty – the intended aim (e.g., learning about objective reality, learning about one's own moral values, etc.) and the choice of a specific strategy to achieve this aim (e.g., expert advice, introspection, etc.). While the first should be tightly associated with metaethical views, associations might be less strong for strategy selection. In Study 2, these two aspects of resolving normative uncertainty were investigated separately. This also allowed for testing whether non-cognitivists intend to reduce their uncertainty by reducing ambivalence. Study 2 also offered a specific strategy for non-naturalists who might have been put off by the naturalist-leaning strategy that had been proposed for realists in Study 1.

Study 2

Study 2 was aimed at replicating and extending the findings from Study 1. Hypotheses for the uncertainty task were the same as in the original study. For the resolving task that addressed the intended aim of resolving uncertainty, it was expected that realists would intend to learn about objective reality, cultural relativists should be interested in learning about their culture or society, individual relativists should be interested in learning about themselves and their own views, and non-cognitivists should be interested in reducing their conflictedness in regard to the respective issue.

For the strategy task that was intended to measure which resolving strategy participants would actually choose, it was expected that realists would either prefer to listen to expert advice from science or philosophy (naturalism) or to their intuition (non-naturalism; Moore, 1903; Ridge, 2019). Cultural relativists were expected to choose to hear opinions from members of their culture. Individual relativists were expected to choose to reflect upon their own views. For non-cognitivists, it was expected that they would indicate that the uncertainty could not be deliberately resolved. The hypotheses for Study 2 had been preregistered on OSF in advance.⁹

Method

Participants

The study was coded in *lab.js* (version 20.2.4; Henninger et al., 2022) and conducted using the *SoSci Survey* software (Leiner, 2022). Participants (N = 200) were recruited via mailing lists from German universities and other German websites that feature online studies. Eight participants were excluded because they indicated not to have responded conscientiously. Further eight participants were excluded, since a relative speed index > 2 indicated that they had rushed the assignment. One of the remaining participants failed more than one attention check, leaving N = 183 participants (129 female, 49 male, 5 non-binary). Participants were aged between 18 and 73 (M = 29.07, SD = 11.44). Most of them had a high educational degree, equivalent to A-level (N = 105). N = 60 had a bachelor's degree or higher. Two further

⁹ Hypotheses regarding post hoc comparisons were only preregistered for associations with the theory task since this task had shown the most consistent associations with other tasks in Study 1: https://osf.io/uxajs

participants had completed an apprenticeship, two were lower secondary school graduates, one had a medium educational level and one was still in school. As reimbursement, participants either received course credit or could take part in a lottery for two 10€ vouchers.

Design and Material

General procedure was similar to Study 1. Some of the presented issues were replaced due to low consensus regarding their classification in the previous study (Table 3). There was one attention check in the metaphor, the theory, the uncertainty, and the category task each. The attention checks appeared at a random location after the first two trials of the respective task.

The order of the tasks was as follows: Metaphor task, disagreement task, theory task, uncertainty task, resolving task, strategy task, and category task. In the metaphor and the theory task, the relativist response option was split in two so that there was one response option for cultural relativism and one for individual relativism in each task. There were no changes to the disagreement or the category task. The full task material can be found in the supplements of this article.

Uncertainty Task. The new response options of the uncertainty task put more emphasis on what the uncertainty ultimately refers to. In the ambivalence response, it was made explicit that the uncertainty does not refer to which answer to the respective issue is correct.

Imagine that someone was uncertain whether **it is morally acceptable for a state to limit immigration**. We will present you with several interpretations of what it may mean to be uncertain about this question.

Which interpretation seems most appropriate to you?

- Ultimately, the person is uncertain which answer to the question is objectively correct.
- (2) Ultimately, the person is uncertain which answer to the question is correct according to the standards of their own culture or society.

- (3) Ultimately, the person is uncertain which answer to the question correctly reflects their own opinion or attitude.
- (4) The uncertainty does not refer to which answer to the question is correct. Instead, it is the result of conflicting feelings, intentions, emotions, or opinions about the question.

Resolving Task. In the novel version of the resolving task, the focus was on the intended aim of resolving uncertainty.

Imagine that you yourself were uncertain whether **it is morally acceptable for a state to limit immigration**. Now you want to become more certain about this question and reduce your uncertainty. We will show you different descriptions of what you would need to do to reduce your uncertainty.

Which description seems most appropriate to you?

- (1) I would have to learn more about objective reality.
- (2) I would have to learn more about my culture or society.
- (3) I would have to learn more about myself and my own views.
- (4) I would have to resolve my conflictedness on this question.

Strategy Task. The strategy task was designed to measure which concrete strategy

participants would deem the most appropriate to reduce their uncertainty.

Imagine that you yourself were uncertain whether **it is morally acceptable for a state to limit immigration**. Now you want to become more certain about this question and reduce your uncertainty. We will show you different strategies you could try to reduce your uncertainty.

Which strategy seems most appropriate to you?

- (1) To listen to the opinion of experts (e.g. from science or philosophy) on this question.
- (2) To listen to my gut feeling or intuition about this question.
- (3) To listen to opinions from society on this question.

(4) To reflect on my own views on this question.

(5) The uncertainty could not be deliberately reduced.

Results

Category Task

Each issue was dominantly categorized as intended, however there was considerable variance regarding the newly introduced issues of meat, slavery, and parents (Table 3). The newly developed issues of birds and color fitted their categories well with > 97% of participants classifying them as intended.

Metaethical Tasks

The distributions of responses to the metaethical tasks across categories and issues can be seen in Figure 6. How participants categorized the presented issues had large effects on how they responded to the metaphor task, $X^2(9) = 1382.37$, p < .001, V(3) = .501, the disagreement task, $X^2(12) = 872.79$, p < .001, V(3) = .396, and the theory task, $X^2(9) = 1551.98$, p < .001, V(3) =.531. Internal consistency of the metaethical tasks was comparable to Study 1. The highest association was found between the metaphor and the theory task, $X^2(9) = 195.92$, p < .001, V(3) = .335 (large effect size), followed by the association between the disagreement and the theory task, $X^2(12) = 154.64$, p < .001, V(3) = .292 (large effect size), and the association between the metaphor and the disagreement task, $X^2 = 94.87$, p < .001, V(3) = .223 (medium effect size). **Uncertainty Task**

The distributions of responses to the uncertainty, the resolving, and the strategy task across categories and issues can be seen in Figure 7. Interpretations of uncertainty varied between categories with a large effect size, $X^2(9) = 1660.23$, p < .001, V(3) = .549. For moral issues, most participants interpreted uncertainty as referring to one's own views (37.7%), followed by interpretations as uncertainty regarding the culturally accepted view (33.8%), uncertainty about which answer is objectively correct (11.8%), and ambivalence (16.6%).

As in Study 1, interpretations of normative uncertainty depended on the participants' metaethical views. Responses to the uncertainty task were associated with responses to the metaphor task, $X^2(9) = 41.20$, p < .001, V(3) = .139 (small effect size), the disagreement task, $X^2 = 44.42$, p < .001, V(3) = .139 (small effect size), and the theory task, $X^2(9) = 93.78$, p < .001, V(3) = .225 (medium effect size). This means that Hypotheses 1a - 1c from the preregistration could be confirmed. As predicted, associations were most consistent in the theory task, where each metaethical view was associated with elevated rates of the corresponding interpretation of normative uncertainty (Figure 8). Hypotheses 4a-4d could therefore also be confirmed. For associations with metaethical views identified via metaphor or disagreement task, only very few post hoc tests were significant.

Resolving Task

What participants intended to achieve by resolving their uncertainty varied between categories with a large effect size, $X^2(9) = 1815.28$, p < .001, V(3) = .574. For moral issues, participants mostly indicated that the aim of resolving uncertainty was to learn about oneself and one's own views (39.7%), followed by learning about one's culture (30.8%), learning about objective reality (17.0%), and reducing one's conflictedness regarding the respective issue (12.5%).

The intended aim was associated with responses to the metaphor task, $X^2(9) = 70.23$, p < .001, V(3) = .192 (medium effect size), responses to the disagreement task, $X^2 = 34.36$, p < .001, V(3) = .116 (small effect size), and responses to the theory task, $X^2(9) = 129.53$, p < .001, V(3) = .269 (medium effect size). This means that Hypotheses 2a-2c from the preregistration could be confirmed. As predicted, each response to the theory task was associated with elevated rates of the respective aim (Figure 9). Hypotheses 5a-5d could therefore be confirmed.

Responses to the resolving task were associated with interpretations of normative uncertainty with a large effect size, $X^2(9) = 176.41$, p < .001, V(3) = .317. Post hoc tests revealed the same consistent pattern as for associations with the theory task. This means that for each

aim of resolving normative uncertainty, there were elevated rates of the respective interpretation of normative uncertainty (Figure 10).

Strategy Task

Responses to the strategy task varied between categories with a large effect size, $X^2(12) = 1601.91$, p < .001, V(3) = .539. For resolving uncertainty about moral issues, participants mostly chose introspection (35.6%) or expert advice (33.3%). This was followed by listening to opinions from society (19.0%) and following one's intuition (8.9%). In 3.2% of the cases, participants indicated that the normative uncertainty could not be deliberately resolved.

Strategy choice depended on the participants' metaethical views. Responses to the strategy task for moral issues were associated with responses to the metaphor task, $X^2 = 43.51$, p < .001, V(3) = .137 (small effect size), responses to the disagreement task, $X^2 = 49.88$, p < .001, V(4) = .123 (small effect size), and responses to the theory task, $X^2 = 70.04$, p < .001, V(3) = .186 (medium effect size). Hypotheses 3a-3c from the preregistration could therefore be confirmed. As predicted, of all metaethical tasks, the theory task showed the most consistent post hoc results with the strategy task. Except for non-cognitivism, each metaethical view was associated with significantly elevated rates of the corresponding resolving strategy (Figure 11). Therefore, Hypotheses 6a-6d, but not Hypothesis 6e, could be confirmed.

Strategy choice was not only associated with metaethical views but also with interpretations of normative uncertainty and with the intended aim of resolving uncertainty. Responses to the strategy task for moral issues depended on responses to the uncertainty task, $X^2 = 85.92$, p < .001, V(3) = .210 (medium effect size), and on responses to the resolving task, $X^2 = 176.54$, p < .001, V(3) = .314 (large effect size). Post hoc tests revealed that the association between strategy choice and uncertainty interpretation was largely driven by relativist views on normative uncertainty (Figure 12). Post hoc tests for associations of strategy choice with the resolving task showed a similar pattern as with the theory task. However, there were no elevated rates of choosing intuition for participants that identified the aim of resolving

uncertainty as learning about objective reality. Yet, participants that indicated that the aim of resolving uncertainty was to reduce one's conflictedness showed elevated rates of indicating that the uncertainty could not be resolved.

Discussion

Study 2 was able to replicate and extend the central findings of the original study while addressing some of its shortcomings. Despite several changes to the study material, the main arguments of this paper were still supported. As in Study 1, realist interpretations of normative uncertainty were the least preferred among lay people, which was explained by underlying metaethical beliefs. Furthermore, by separating the intended aim from concrete strategy selection, Study 2 was able to extend findings on how people think about resolving normative uncertainty.

Adaptations to Metaphor and Theory Task

One concern regarding the original study was the inconsistency of findings in regard to relativism. This was addressed by adapting the metaphor and the theory task so that they would differentiate between cultural and individual relativism. The results of Study 2 largely supported these adaptations. As predicted, cultural and individual relativism as identified by the theory task were now consistently linked to the respective responses to the uncertainty, the resolving, and the strategy task. For the metaphor task, associations of cultural and individual relativism could only be shown with the resolving task but not with the uncertainty or the strategy task. Overall, these findings showed that the distinction between cultural and individual relativism is necessary for understanding lay views on normative uncertainty and that the adapted metaphor task and especially the theory task were able to capture this differentiation.

Adaptations to the Uncertainty Task

The second limitation of the original study was the high number of cases where normative uncertainty was interpreted as ambivalence. This concern was addressed by adapting the ambivalence response option for the uncertainty task so that it would only be attractive to non-cognitivist but not to cognitivist participants. The overall rates of interpreting uncertainty as ambivalence in Study 2 were almost halved in comparison to Study 1, all while preserving the link to non-cognitivism. This reduction of ambivalence interpretations could also not be attributed to differences in the rates of non-cognitivist respondents, since these rates were stable across Study 1 and 2. Therefore, the adaptation of the non-cognitivist response option successfully sharpened the contrast between cognitivist and non-cognitivist interpretations of normative uncertainty.

A further concern regarding Study 1 was that some realists might have been drawn to cultural or individual relativist interpretations of uncertainty, despite being ultimately uncertain about objective moral facts. It was expected that changing the cognitivist response options in the uncertainty task, so that they would stress the aspect of what the uncertainty ultimately refers to, would reduce such effects. Yet, Study 2 did not support this argument. The overall effect sizes of associations between the uncertainty task and the metaethical tasks were smaller than in the original study, even after accounting for the change in degrees of freedom. Similarly, associations between realism and realist interpretations of normative uncertainty were less pronounced in Study 2 than in Study 1. Reexamining findings from Study 1 revealed that associations between the uncertainty and the metaethical tasks in regard to realism had already been of similar magnitude as associations between some of the metaethical tasks themselves. For example, the association between realist responses to the uncertainty and the theory task were similar to those between realist responses to the metaphor and the theory task. This indicates that the realist response option in the original uncertainty task already exhibited high validity and was well equipped to reveal associations with metaethical views. Concerns about realists systematically misinterpreting normative uncertainty in relativist ways may therefore have been unwarranted. Overall, the cognitivist response options in the uncertainty task in Study 1 were better suited than those in Study 2 to identify associations between uncertainty interpretations and metaethical views.

Resolving and Strategy Task

The most substantial novelty of Study 2 concerned the resolving of normative uncertainty. Both the resolving and the strategy task showed associations with the respective metaethical positions. As expected, associations of the strategy task were less pronounced than those of the resolving task. Associations between the strategy and the resolving task were strong and consistent.

In the strategy task, there was some support for the argument from Eriksson and Francén Olinder (2016) according to which individual relativists would turn towards outside factors to resolve their normative uncertainty, as a substantial share of individual relativists actually selected expert advice. Yet, the large majority of individual relativists thought that introspection would be the most appropriate resolving strategy.

It was also argued that realists might try to resolve their uncertainty by referring to consensus information or by focusing on their own mental states (Ayars & Nichols, 2020; MacAskill et al., 2020). For the former, there was no empirical support in Study 2. However, there were some realists that chose introspection and there was a significant association between realism and choosing intuition. These findings indicate that realist lay people do not necessarily refer to external sources to resolve their normative uncertainty.

The association between realism and intuition also suggested that non-naturalism might be a relevant factor in lay people's metaethical views. Even though the overall rates of choosing intuition were rather low in the current study, future studies on lay metaethics might benefit from distinguishing between naturalist and non-naturalist views.

Further Limitations

The presented studies largely built on metaethical measurements described in Wagner et al. (2021). Of all the metaethical tasks, the theory task consistently showed the strongest associations with the uncertainty-related tasks. One reason might be that the theory task was the most explicit measure of metaethical intuition in these studies. Therefore, it might have had higher validity than the other measurements, such as the metaphor task where some participants stated that they had problems understanding it. In future studies, the effects of implicit, semi-implicit and explicit measures should be investigated more systematically. First steps in this direction have been taken by Pölzler and Wright (2020), and Wagner et al. (2021).

The current findings also raised some worries in regard to the disagreement task. In both studies, relativism as identified by metaphor or theory task was not associated with relativist but with non-cognitivist responses to the disagreement task (see Supplementary Tables). One explanation for this might be that the structure of the disagreement task increases the chance of measurement error since the identification of a metaethical view requires responses to two versions of the task (Pölzler & Wright, 2020). However, this would not explain why relativist participants chose non-cognitivist interpretations instead. As disagreement tasks are one of the most used types of metaethical measurements in the literature (e.g., Goodwin & Darley, 2008, 2012; Sarkissian et al., 2011; Wright & Pölzler, 2021), a further validation of the disagreement task is required.

Another concern regarding the use of disagreement tasks was raised by Bush and Moss (2020). They argue that participants might misinterpret the task material as disagreement about non-moral aspects of the described issues, instead of genuine moral disagreement (see also Rowland, 2021). Similarly, uncertainty might have been misinterpreted as uncertainty regarding non-normative facts, such as what the consequences of limiting immigration would be, and not as normative uncertainty. If participants interpreted uncertainty and disagreement as referring to non-normative aspects, this would artificially increase the proportion of realist responses to these tasks, since these non-normative aspects would be construed by many as objective facts. Even though this would explain the association between realists in the disagreement task and the responses to the uncertainty-related tasks, this would not affect associations of those latter tasks with other metaethical measurements. Therefore, the main findings of the presented studies seem to be largely robust against this criticism. Lastly, the generalizability of the present findings needs to be discussed since both studies largely relied on convenience samples from a university context. Participants were thus expected to be younger and more educated than a broader public. There is some limited data suggesting age effects on metaethical views and experiences of normative uncertainty (Beebe & Sackris, 2016; Jabarian, 2020). However, these findings do neither refer to interpretations of normative uncertainty nor to associations between uncertainty interpretations and metaethical views. Regarding education, it seems plausible that a lack of experience with philosophical concepts would increase measurement error or introduce bias. The use of comprehension checks can help to identify such effects. However, excluding participants that fail these checks might further increase the selectivity of the final sample (see for example Pölzler & Wright, 2020). One of the central challenges of experimental philosophy remains the operationalization of abstract and complex constructs in a way that makes it possible to analyze them in a philosophically naïve population. The empirical investigation of metaethics is no exception to this.

General Discussion

The empirical investigation of normative uncertainty is still in its infancy. In this paper, I offer a basis for such research by investigating lay concepts of normative uncertainty and how they depend on metaethical assumptions.

The current studies revealed a fundamental discrepancy between the way lay people think about normative uncertainty and how normative uncertainty has been construed in the uncertaintism discourse. This also affected self-reports of how people would try to resolve their normative uncertainty. Yet, besides metaethical intuitions, there are other factors that might affect lay views on this topic, such as characteristics of the moral issue itself. For example, moral issues might differ in the salience of the underlying value conflict. In issues where this conflict is very obvious, such as trolley-style cases, interpretations of normative uncertainty as ambivalence might be more prevalent than in issues where the identification of the underlying value conflict requires deeper analysis. Further, uncertainty even about the same moral issue might be interpreted differently by the same person, depending on which aspect of the issue the uncertainty refers to. This can be illustrated for normative uncertainty regarding the permissibility of abortion. A realist could be uncertain regarding this issue because they do not know which notion of identity is best suited to inform us about the moral status of the fetus (Moller, 2011). This would be normative uncertainty regarding, according to the realist, objective facts. However, the same realist could also be uncertain because they are torn between supporting the rights of the pregnant person and those of the fetus, which might lead to interpretations of ambivalence rather than uncertainty about objective facts. Research has already started to address which aspects of a moral issue affect its metaethical evaluation (Goodwin & Darley, 2012; Wagner et al., 2021). Similarly, situational factors that affect the interpretation of normative uncertainty should receive further attention.

It should also be discussed to which degree the current findings extend beyond selfreports and inform us about the participants' actual moral practice (Colebrook, 2021). There are multiple ways in which moral practice might differ from reported metaethical views: For example, despite indicating a non-cognitivist view, a participant might have moral judgments that show characteristics of beliefs (see Cohen et al., 2022), their neural activity during a moral evaluation might more closely resemble reasoning about facts than about preferences (see Theriault et al., 2017, 2020), or their reasoning under normative uncertainty might more closely refer to ignorance of a fact than to ambivalence. If explicit metaethical beliefs and actual moral practice were completely unrelated, the current findings would only show that lay peoples' beliefs regarding metaethics and normative uncertainty are somewhat consistent. On the other hand, if explicit beliefs were accurate reflections of the participant's moral practice, the current findings would suggest that there are intra- and interindividual differences in how lay people reason under normative uncertainty, depending on their metaethical views about the respective issue. A moderate interpretation of the current results might acknowledge that explicit metaethical beliefs do not equal moral practice but that it is likely that there is at least some association between them. Such an association might be explained by introspective access to characteristics of one's own moral judgments that might inform lay people's metaethical beliefs.

The current findings are inconclusive in this regard. On one hand, responses to the less explicit disagreement tasks showed lower associations with uncertainty-related measures than the more explicit tasks, which would speak for a discrepancy between explicit beliefs and moral practice. On the other hand, associations between metaethical views and the strategy task, which has a more pronounced behavioral aspect than the other tasks, indicate that a moderate interpretation of the findings might be warranted. Further research should address the relationship between beliefs and moral practice, including reasoning under normative uncertainty. Recently, behavioral and neuroimaging paradigms have been proposed that might be suitable for this task (e.g., Cohen et al., 2022; Theriault et al., 2017, 2020).

Implications for Empirical Research on Normative Uncertainty

So far, only a few researchers have tried to investigate normative uncertainty by empirical means (Costa-Gomes & Schönegger, 2022; Jabarian, 2020). However, it can be expected that more researchers will take on this task in the near future, given its relevance for moral decision-making and the increasing body of respective philosophical literature. In other experimental philosophy work, researchers have drawn from thought experiments to generate their stimulus material (Feltz, 2009). Such thought experiments also exist in the normative uncertainty literature (e.g., Harman, 2015; Lockhart, 2000; Moller, 2011). Some of these utilize terms such as *subjective probabilities* or *correct moral theory* to describe the parameters of normative uncertainty. The current findings suggest that this terminology might not be suited to describe normative uncertainty for a lay public. Describing normative uncertainty as subjective probabilities or as the strength of a belief that some moral view is correct might lead to antirealist participants interpreting this uncertainty as descriptive rather than normative, or might even be nonsensical to non-cognitivist participants. The current findings should encourage
researchers to consider the pluralism in lay people's metaethical intuitions either by employing a language that is agnostic towards metaethical presumptions or by explicitly addressing the here described variants of normative uncertainty.

Conclusion

Validly measuring folk opinions on metaethics is far from trivial and some have argued that such attempts are bound to fail (Bush & Moss, 2020; Colebrook, 2021). However, even if participants may not be realist or anti-realist to the proportions that are described here, there are still interesting correlations between their responses to the metaethical tasks and their interpretations of normative uncertainty. At the very least, the current studies show that a substantial share of participants does not interpret normative uncertainty in the way it is construed in most of the uncertaintism literature. This calls for consideration in the development of any empirical studies on choice under normative uncertainty.

References

Agresti, A. (2007). An introduction to categorical data analysis. Wiley.

- Alsaad, A. K. (2021). Ethical judgment, subjective norms, and ethical consumption: The moderating role of moral certainty. *Journal of Retailing and Consumer Services*, *59*, 102380. https://doi.org/10.1016/j.jretconser.2020.102380
- Alsaad, A. K., Elrehail, H., & Saif-Alyousfi, A. Y. H. (2021). The interaction among religiosity, moral intensity and moral certainty in predicting ethical consumption: A study of Muslim consumers. *International Journal of Consumer Studies*. Advance online publication. https://doi.org/10.1111/ijcs.12688
- Ayars, A., & Nichols, S. (2020). Rational learners and metaethics: Universalism, relativism, and evidence from consensus. *Mind & Language*, *35*(1), 67–89. https://doi.org/10.1111/mila.12232
- Beddor, B. (2020). Fallibility for expressivists. *Australasian Journal of Philosophy*, *98*(4), 763–777. https://doi.org/10.1080/00048402.2019.1699586
- Beebe, J., Qiaoan, R., Wysocki, T., & Endara, M. A. (2015). Moral objectivism in cross-cultural perspective. *Journal of Cognition and Culture*, 15(3-4), 386–401. https://doi.org/10.1163/15685373-12342157
- Beebe, J., & Sackris, D. (2016). Moral objectivism across the lifespan. *Philosophical Psychology*, *29*(6), 912–929. https://doi.org/10.1080/09515089.2016.1174843
- Bergsma, W. (2013). A bias-correction for Cramér's V and Tschuprow's T. *Journal of the Korean Statistical Society*, *42*(3), 323–328. https://doi.org/10.1016/j.jkss.2012.10.002
- Björkholm, S., Bykvist, K., & Olson, J. (2021). Quasi-realism and normative certitude. *Synthese*, *198*(8), 7861–7869. https://doi.org/10.1007/s11229-020-02553-y
- Blackburn, S. (1993). Essays in quasi-realism. Oxford University Press.
- Brink, D. O. (1989). *Moral realism and the foundation of ethics. Cambridge studies in philosophy*. Cambridge University Press.

Bush, L. S., & Moss, D. (2020). Misunderstanding metaethics: Difficulties measuring folk objectivism and relativism. *Diametros*, 17(64), 6–21. https://doi.org/10.33392/diam.1495

Bykvist, K. (2017). Moral uncertainty. *Philosophy Compass*, *12*(3), e12408. https://doi.org/10.1111/phc3.12408

Bykvist, K., & Olson, J. (2017). Non-cognitivism and fundamental moral certitude: Reply to Eriksson and Francén Olinder. *Australasian Journal of Philosophy*, 95(4), 794–799. https://doi.org/10.1080/00048402.2016.1269352

Cohen, H., Nissan-Rozen, I., & Maril, A. (2022). Empirical evidence for moral bayesianism. *Philosophical Psychology*, Advance online publication.

https://doi.org/10.1080/09515089.2022.2096430

Cohen, J. (1988). Statistical power analysis for the behavioral sciences. Routledge.

- Colebrook, R. (2021). The irrationality of folk metaethics. *Philosophical Psychology*, *34*(5), 684–720. https://doi.org/10.1080/09515089.2021.1915970
- Costa-Gomes, M., & Schönegger, P. (2022). *Charitable giving under normative uncertainty: Experimental evidence on the behavioral impact of normative expert advice* [Manuscript in preparation].
- Egan, A. (2007). Quasi-realism and fundamental moral error. *Australasian Journal of Philosophy*, *85*(2), 205–219. https://doi.org/10.1080/00048400701342988

Eriksson, J., & Francén Olinder, R. (2016). Non-cognitivism and the classification account of moral uncertainty. *Australasian Journal of Philosophy*, 94(4), 719–735. https://doi.org/10.1080/00048402.2016.1149199

- Ewing, A. C. (1948). The definition of good. Routledge.
- Feltz, A. (2009). Experimental philosophy. Analyse & Kritik (2), 201–219.
- Foot, P. (1967). The problem of abortion and the doctrine of the double effect. *Oxford Review*, *5*, 5–15.

Gibbard, A. (1990). Wise choices, apt feelings. Harvard University Press.

- Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, *106*(3), 1339–1366. https://doi.org/10.1016/j.cognition.2007.06.007
- Goodwin, G. P., & Darley, J. M. (2012). Why are some moral beliefs perceived to be more objective than others? *Journal of Experimental Social Psychology*, *48*(1), 250–256. https://doi.org/10.1016/j.jesp.2011.08.006
- Gustafsson, J. E. (2022). Second thoughts about My Favourite Theory. *Pacific Philosophical Quarterly*, Article papq.12408. Advance online publication.

https://doi.org/10.1111/papq.12408

- Gustafsson, J. E., & Torpman, O. (2014). In defence of My Favourite Theory. *Pacific Philosophical Quarterly*, 95(2), 159–174. https://doi.org/10.1111/papq.12022
- Guzmán, R. A., Barbato, M. T., Sznycer, D., & Cosmides, L. (2022). A moral trade-off system produces intuitive judgments that are rational and coherent and strike a balance between conflicting moral values. *Proceedings of the National Academy of Sciences of the United States of America*, *119*(42), e2214005119.

https://doi.org/10.1073/pnas.2214005119

- Harman, E. (2015). The irrelevance of moral uncertainty. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 10, pp. 53–79). Oxford University Press.
- Henninger, F., Shevchenko, Y., Mertens, U., Kieslich, P. J., & Hilbig, B. E. (2022). *lab.js: A free, open, online experiment builder* [Computer software]. Zenodo.

Hope, A. C. A. (1968). A simplified Monte Carlo significance test procedure. *Journal of the Royal Statistical Society: Series B (Methodological)*, *30*(3), 582–598.
https://doi.org/10.1111/j.2517-6161.1968.tb00759.x

Jabarian, B. (2020, May 30). *Moral uncertainty, social values and economic preferences* [Paper presentation]. Princeton University, Princeton, NJ, USA.

Jafari, M., & Ansari-Pour, N. (2019). Why, when and how to adjust your p values? *Cell Journal*,

20(4), 604-607. https://doi.org/10.22074/cellj.2019.5992

Joyce, R. (2011). Moral fictionalism. Philosophy Now, 82, 14–17.

- Köhler, S. (2015). What is the problem with fundamental moral error? *Australasian Journal of Philosophy*, *93*(1), 161–165. https://doi.org/10.1080/00048402.2014.928736
- Lam, G. (2020). Making quasi-realists admit of fundamental moral fallibility. *Australasian Journal of Philosophy*, *98*(2), 294–303. https://doi.org/10.1080/00048402.2019.1611892
- Leiner, D. J. (2019). Too fast, too straight, too weird: Non-reactive indicators for meaningless data in internet surveys. *Survey Research Methods*, *13*(3), 229–248. https://doi.org/10.18148/srm/2019.v13i3.7403
- Leiner, D. J. (2022). *SoSci Survey* (Version 3.2.55) [Computer software]. https://www.soscisurvey.de
- Lenman, J. (2003). *Non-cognitivism and the dimensions of evaluative judgment*. Brown Electronic Article Review Service.

https://www.brown.edu/Departments/Philosophy/bears/0301lenm.html

Lockhart, T. (2000). Moral uncertainty and its consequences. Oxford University Press.

Lutz, M., & Lenman, J. (2021). Moral naturalism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy.* https://plato.stanford.edu/archives/spr2021/entries/naturalism-moral/

MacAskill, W., Bykvist, K., & Ord, T. (Eds.). (2020). *Moral uncertainty*. Oxford University Press.

Mackie, J. L. (1977). *Ethics: Inventing right and wrong*. Penguin Books.

- Makins, N. (2021). Attitudinal ambivalence: Moral uncertainty for non-cognitivists. *Australasian Journal of Philosophy*, 1–15. https://doi.org/10.1080/00048402.2021.1908380
- Mata, A. (2019). Social metacognition in moral judgment: Decisional conflict promotes perspective taking. *Journal of Personality and Social Psychology*, *117*(6), 1061–1082. https://doi.org/10.1037/pspa0000170

Moller, D. (2011). Abortion and moral risk. *Philosophy*, 86(3), 425–443.

https://doi.org/10.1017/S0031819111000222

Moore, G. E. (1903). Principia ethica. Cambridge University Press.

- Nissan-Rozen, I. (2015). Against moral hedging. *Economics and Philosophy*, *31*(3), 349–369. https://doi.org/10.1017/S0266267115000206
- Pendlebury, M. (2011). Objectivism versus realism. *The Philosophical Forum*, *42*(1), 79–104. https://doi.org/10.1111/j.1467-9191.2010.00379.x
- Petrocelli, J. V., Tormala, Z. L., & Rucker, D. D. (2007). Unpacking attitude certainty: Attitude clarity and attitude correctness. *Journal of Personality and Social Psychology*, *92*(1), 30–41. https://doi.org/10.1037/0022-3514.92.1.30
- Pölzler, T., & Wright, J. C. (2019). Empirical research on folk moral objectivism. *Philosophy Compass*, 14(5), e12589. https://doi.org/10.1111/phc3.12589
- Pölzler, T., & Wright, J. C. (2020). Anti-realist pluralism: A new approach to folk metaethics. *Review of Philosophy and Psychology*, *11*(1), 53–82. https://doi.org/10.1007/s13164-019-00447-8
- Railton, P. (1986). Moral realism. *The Philosophical Review*, *95*(2), 163–207. https://doi.org/10.2307/2185589
- Ridge, M. (2019). Moral non-naturalism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy.* https://plato.stanford.edu/archives/fall2019/entries/moral-non-naturalism/
- Robinson, P., & Steele, K. (2022). Moral uncertainty, noncognitivism, and the multi-objective story. *Noûs*, Article nous.12435. Advance online publication.

https://doi.org/10.1111/nous.12435

Rosenthal, C. (2021). What decision theory can't tell us about moral uncertainty. *Philosophical Studies*, *178*(10), 3085–3105. https://doi.org/10.1007/s11098-020-01571-3

Rowland, R. (2021). Moral disagreement. New problems of philosophy. Routledge.

- Sarkissian, H., Park, J., Tien, D., Wright, J. C., & Knobe, J. (2011). Folk moral relativism. *Mind & Language*, *26*(4), 482–505. https://doi.org/10.1111/j.1468-0017.2011.01428.x
- Sepielli, A. (2012). Normative uncertainty for non-cognitivists. *Philosophical Studies*, *160*(2), 191–207. https://doi.org/10.1007/s11098-011-9713-6
- Sepielli, A. (2017). How moral uncertaintism can be both true and interesting. In M. C. Timmons (Ed.), Oxford Studies in Normative Ethics (pp. 98–116). Oxford University Press. https://doi.org/10.1093/oso/9780198808930.003.0006

Shafer-Landau, R. (2003). Moral realism: A defence. Clarendon.

- Sharpe, D. (2015). Chi-square test is statistically significant: Now what? *Practical Assessment, Research, and Evaluation, 20*(8). https://doi.org/10.7275/TBFA-X148
- Shore, B. (1990). Human ambivalence and the structuring of moral values. *Ethos*, 18(2), 165–179.

Skitka, L. J., Hanson, B. E., Morgan, G. S., & Wisneski, D. C. (2021). The psychology of moral conviction. *Annual Review of Psychology*, 72, 347–366. https://doi.org/10.1146/annurev-psych-063020-030612

- Smith, M. (2002). Evaluation, uncertainty and motivation. *Ethical Theory and Moral Practice*, 5(3), 305–320. https://doi.org/10.1023/A:1019675327207
- Staffel, J. (2019). Expressivism, normative uncertainty, and arguments for probabilism. In T. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology* (161-189). Oxford University Press.
- Tarsney, C. (2021). Normative externalism by Brian Weatherson. *Mind*, *130*(519), 1018–1028. https://doi.org/10.1093/mind/fzaa018
- Theriault, J., Waytz, A., Heiphetz, L., & Young, L. (2017). Examining overlap in behavioral and neural representations of morals, facts, and preferences. *Journal of Experimental Psychology. General, 146*(11), 1586–1605. https://doi.org/10.1037/xge0000350

Theriault, J., Waytz, A., Heiphetz, L., & Young, L. (2020). Theory of mind network activity is associated with metaethical judgment. An item analysis. *Neuropsychologia*, *143*, 107475. https://doi.org/10.1016/j.neuropsychologia.2020.107475

Tilley, J. J. (2000). Cultural relativism. *Human Rights Quarterly*, 22(2), 501–547.

Turp, M.-J. (2018). Normativity, realism and emotional experience. *Philosophia*. Advance online publication. https://doi.org/10.1007/s11406-018-9984-7

van Harreveld, F., van der Pligt, J., & Liver, Y. N. de (2009). The agony of ambivalence and ways to resolve it: Introducing the MAID model. *Personality and Social Psychology Review*, *13*(1), 45–61. https://doi.org/10.1177/1088868308324518

- van Roojen, M. (2018). Moral cognitivism vs. non-cognitivism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/fall2018/entries/moral-cognitivism/
- Wagner, J. M. J., Pölzler, T., & Wright, J. C. (2021). Implicit metaethical intuitions: Validating and employing a new IAT procedure. *Review of Philosophy and Psychology*. Advance online publication. https://doi.org/10.1007/s13164-021-00572-3

Weatherson, B. (2019). Normative externalism. 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 & Social Psychology Bulletin*, 34(11), 1461–1476.
https://doi.org/10.1177/0146167208322557

Wright, J. C., & Pölzler, T. (2021). Should morality be abolished? An empirical challenge to the argument from intolerance. *Philosophical Psychology*, 1–36. https://doi.org/10.1080/09515089.2021.1983160

Zijlstra, L. (2019). Folk moral objectivism and its measurement. *Journal of Experimental Social Psychology*, *84*, 103807. https://doi.org/10.1016/j.jesp.2019.04.005 Zijlstra, L. (2021). Are people implicitly moral objectivists? *Review of Philosophy and Psychology*.

Advance online publication. https://doi.org/10.1007/s13164-021-00593-y

Table 1

Overview of Metaethical Notions and the Assumed Psychological States of Normative

Uncertainty

		Cognitivism		Non-cognitivism
Criterion	Realism	Cultural relativism	Individual relativism	
Understanding of moral judgments	Beliefs about universal, objective facts	Beliefs about facts that can vary between cultures	Beliefs about facts that can vary within cultures	Non-cognitive states (desires, intentions,)
Interpretation of normative uncertainty	Uncertainty regarding universal, objective moral facts	Uncertainty regarding culturally accepted moral values	Uncertainty regarding own moral values	Ambivalent non-cognitive states

Table 2

lssue	Question	Intended category	Science	Morality	Social conven- tion	Personal prefe- rence
Brain	Do brain scans represent a reliable scientific method?	Science	96.15%	0.00%	2.20%	1.65%
Sun	Does the earth revolve around the sun?		98.90%	0.00%	1.10%	0.00%
Immigration	Is it morally acceptable for a state to limit immigration?	Morality	3.85%	87.36%	6.59%	2.20%
Needy	Do we have a moral duty to help those in need?		1.10%	90.66%	4.95%	3.30%
Rape	ls it wrong to rape someone?		12.64%	76.92%	9.89%	0.55%
Mouthful	Should one speak with a full mouth?	Social convention	0.00%	1.65%	95.60%	2.75%
Pyjamas	Is it wrong to come to seminar in your pyjamas?		0.00%	0.55%	84.62%	14.84%
Mozart	Was Mozart a better musician than Lady Gaga?	Personal taste	1.65%	0.00%	3.85%	94.51%
Sandals	Do socks look terrible in sandals?		2.20%	1.65%	23.08%	73.08%

Description of Issues and How These Were Categorized by Participants in Study 1

Table 3

					Social	Personal
lecuo	Question	Intended	Science	Morality	conven-	prefe-
	Question	category		woranty		Tence
Birds	Are there more birds than frogs?	Science	97.27%	0.00%	1.09%	1.64%
Sun	Does the earth revolve around the sun?		99.45%	0.00%	0.55%	0.00%
Immigration	Is it morally acceptable for a state to limit immigration?	Morality	2.73%	83.06%	11.48%	2.73%
Needy	Do we have a moral duty to help those in need?		0.55%	89.62%	7.10%	2.73%
Meat	ls it reprehensibl e to eat meat?		6.56%	46.45%	7.65%	39.34%
Slavery	ls human trafficking wrong?		7.65%	78.69%	12.02%	1.64%
Mouthful	Should one speak with a full mouth?	Social convention	0.55%	4.92%	87.43%	7.10%
Parents	Would it be wrong to address one's parents informally?		0.00%	2.73%	77.05%	20.22%
Mozart	Was Mozart a better musician than Lady Gaga?	Personal taste	4.37%	0.00%	2.73%	92.9%
Color	ls yellow a more beautiful		0.55%	0.00%	0.55%	98.91%

Description of Issues and How These Were Categorized by Participants in Study 2

		Intended			Social conven-	Personal prefe-
Issue	Question	category	Science	Morality	tion	rence
	color than blue?					

Distributions of Responses to the Metaethical Tasks Across Issues and Categories in Study 1



Disagreement







Distributions of Responses to the Uncertainty and the Resolving Task Across Issues and

Categories in Study 1



Resolving





Post Hoc Plot for Associations of Metaethical Tasks With the Uncertainty Task in Study 1



were more (less) people in the respective cell than what could have been expected from the marginal distributions. Adjusted standardized residuals can be standardized residuals for each combination of responses to the metaethical tasks and the uncertainty task. Positive (negative) values indicate that there Note. Upper row shows distributions of responses to the uncertainty task for each response to the metaethical tasks. Lower row shows adjusted compared across different Chi squared tests.

 $p < .05, *^{*}p < .01, *^{**}p < .001.$



Post Hoc Plot for Associations of the Uncertainty Task With the Resolving Task in Study 1



Distributions of Responses to the Metaethical Tasks Across Issues and Categories in Study 2











Distributions of Responses to the Uncertainty, the Resolving, and the Strategy Task Across Issues



Uncertainty



Resolving









Post Hoc Plot for Associations of Metaethical Tasks With the Uncertainty Task in Study 2



ი
Ð
F
50
:

Post Hoc Plot for Associations of Metaethical Tasks With the Resolving Task in Study 2



Post Hoc Plot for Associations of the Uncertainty Task With the Resolving Task in Study 2



-
• •
a)
Ξ.
=
-
b0
. <u> </u>

Post Hoc Plot for Associations of Metaethical Tasks With the Strategy Task in Study 2



Post Hoc Plot for Associations of Uncertainty and Resolving Task With the Strategy Task in Study 2



Appendix

Table S1

Adjusted Standardized Residuals for Associations of Participants' Categorization of Issues With

All Other Tasks in Study 1

	Category				
– Response	Science	Morality	Social convention	Personal preference	
Metaphor					
Realist	17.44***	4.81***	-11.94***	-10.75***	
Relativist	-9.36***	0.50	13.33***	-5.01***	
Non- cognitivist	-8.48***	-5.99***	-2.41	18.01***	
Disagreement					
Realist	18.32***	6.68***	-12.21***	-13.43***	
Cultural relativist	-3.31*	-2.57	8.28***	-2.54	
Individual relativist	-4.07***	0.05	0.81	3.32*	
Non- cognitivist	-10.70***	-4.73***	4.01**	12.08***	
Non-standard	-4.20***	-0.83	4.58***	0.42	
Theory					
Realist	25.63***	3.07*	-14.58***	-14.54***	
Relativist	-10.70***	5.11***	12.15***	-7.45***	
Non- cognitivist	-14.98***	-7.99***	2.64	21.60***	
Uncertainty					
Ignorance of objective fact	29.08***	-4.30***	-13.23***	-11.45***	
Ignorance of culturally accepted view	-10.40***	-2.57	19.74***	-7.37***	
lgnorance of own view	-9.02***	1.26	-2.55	10.73***	
Ambivalence	-10.98***	5.89***	-4.06***	9.27***	

	Category			
Response	Science	Morality	Social convention	Personal preference
Resolving				
Look for answer in science or philosophy	29.38***	0.34	-15.60***	-13.95***
Look for answer in own culture	-13.13***	-6.52***	25.80***	-7.25***
Increase clarity regarding own view	-14.86***	6.54***	-7.96***	17.08***
Not purposefully resolvable	-4.93***	-1.43	-1.80	8.32***

Note. Positive values indicate a higher frequency than would have been expected without an association of the respective variables. Negative values indicate a lower frequency than would have been expected without an association.

 $p^* < .05, p^{**} < .01, p^{***} < .001.$

Table S2

Adjusted Standardized Residuals for Associations Between Metaethical Tasks for Issues That

		Theory response	
Response	Realist	Relativist	Non-cognitivist
Metaphor			
Realist	8.15***	-3.71**	-5.37***
Relativist	-6.19***	5.11***	1.28
Non-cognitivist	-2.80*	-2.13	5.98***
Disagreement			
Realist	12.61***	-5.79***	-8.23***
Cultural relativist	-0.85	0.58	0.33
Individual relativist	-4.18***	2.08	2.53
Non-cognitivist	-8.08***	3.19*	5.90***
Non-standard	-3.70**	2.12	1.91
		Metaphor response	
Response	Realist	Relativist	Non-cognitivist
Disagreement			
Realist	8.64***	-5.43***	-4.64***
Cultural relativist	-0.56	0.15	0.59
Individual relativist	-1.79	0.12	2.46
Non-cognitivist	-6.76***	5.03***	2.47
Non-standard	-1.95	1.24	1.03

Were Classified as Moral by the Participants in Study 1

Note. Positive values indicate a higher frequency than would have been expected without an association of the respective variables. Negative values indicate a lower frequency than would have been expected without an association.

 $p^* < .05, p^{**} < .01, p^{***} < .001.$

Table S3

Adjusted Standardized Residuals for Associations of Participants' Categorization of Issues With

		Cate	egory	
– Response	Science	Morality	Social convention	Personal preference
Metaphor				
Realist	28.69***	-4.31***	-10.01***	-13.18***
Cultural relativist	-11.35***	8.35***	17.08***	-13.73***
Individual relativist	-12.17***	-2.21	-3.85**	17.21***
Non- cognitivist	-5.11***	-2.85	-5.20***	12.52***
Disagreement				
Realist	23.41***	2.53	-10.35***	-15.05***
Cultural relativist	-3.95**	-1.05	11.34***	-5.61***
Individual relativist	-6.43***	-0.72	0.66	6.16***
Non- cognitivist	-10.92***	-1.46	-2.02	13.59***
Non-standard	-7.04***	0.12	7.28***	-0.21
Theory				
Realist	31.55***	-4.25***	-11.26***	-14.78***
Cultural relativist	-11.97***	7.62***	16.04***	-11.43***
Individual relativist	-8.33***	3.93**	0.41	3.33*
Non- cognitivist	-12.29***	-6.58***	-5.12***	23.05***

All Other Tasks in Study 2

	Category				
Response	Science	Morality	Social convention	Personal preference	
Uncertainty					
Objectively correct answer	32.93***	-8.57***	-10.60***	-12.17***	
Culturally accepted answer	-11.52***	4.58***	19.64***	-11.99***	
Own opinion	-11.29***	8.33***	-3.21*	4.83***	
Ambivalence	-10.03***	-4.56***	-6.40***	20.00***	
Resolving					
Learn about objective reality	34.19***	-6.76***	-12.25***	-13.73***	
Learn about one's culture or society	-12.40***	2.89	22.93***	-12.42***	
Learn about oneself and one's opinions	-14.37***	6.01***	-5.66***	12.37***	
Reduce conflictedness	-8.85***	-2.74	-5.50***	16.18***	
Strategy					
Expert advice	29.77***	-0.49	-12.11***	-16.24***	
Intuition	-7.66***	-2.94	-1.21	11.34***	
Opinions from society	-11.19***	-0.43	23.05***	-10.21***	
Introspection	-12.92***	5.23***	-4.79***	11.03***	
Not purposefully resolvable	-5.59***	-3.87**	-3.51**	12.47***	

Note. Positive values indicate a higher frequency than would have been expected without an association of the respective variables. Negative values indicate a lower frequency than would have been expected without an association.

 $^{*}p < .05, \, ^{**}p < .01, \, ^{***}p < .001.$

Table S4

Individual

relativist

Non-cognitivist

Non-standard

-2.35

-5.39***

-0.35

Adjusted Standardized Residuals for Associations Between Metaethical Tasks for Issues That

	Theory response				
			Individual		
Response	Realist	Cultural relativist	relativist	Non-cognitivist	
Metaphor					
Realist	9.92***	-3.80**	-3.48**	-1.82	
Cultural relativist	-2.83	8.41***	-5.35***	-1.98	
Individual relativist	-4.05***	-4.13***	7.60***	1.33	
Non-cognitivist	-2.40	-2.91	2.37	3.65**	
Disagreement					
Realist	10.84***	-0.56	-4.63***	-5.65***	
Cultural relativist	-1.50	2.91	-1.27	-0.77	
Individual relativist	-3.22*	-0.42	2.11	1.63	
Non-cognitivist	-6.45***	-1.38	3.78**	4.43***	
Non-standard	-2.14	0.94	0.22	0.81	
		Metaphor re	esponse		
Response	Realist	Cultural relativist	Individual relativist	Non-cognitivist	
Disagreement					
Realist	7.82***	0.37	-4.94***	-3.63**	
Cultural relativist	-1.33	1.82	-0.91	0.05	

Were Classified as Moral by the Participants in Study 2

Note. Positive values indicate a higher frequency than would have been expected without an association of the respective variables. Negative values indicate a lower frequency than would have been expected without an association.

0.28

-1.24

-0.39

2.28

3.94**

0.02

-0.68

3.32*

1.05

 $^{*}p < .05, \, ^{**}p < .01, \, ^{***}p < .001.$

Task Material

Study 1

Metaphor Task.

Is it morally acceptable for a state to limit immigration? Here, we are not interested in what is the correct answer to this question. Rather, we will present you with metaphors about whether there is a correct answer to the question, and if yes, what it is that makes this answer correct.

Which of these metaphors seems most appropriate to you?

- (1) There is a correct answer to this question. It is "out there in the world". [REALISM]
- (2) There is a correct answer to this question. It is "invented or created by individuals or societies". [RELATIVISM]
- (3) There is not a correct answer to this question. Any response to this question is just like someone saying "Boo!" or "Hurray!" [NON-COGNITIVISM]

Disagreement Task.

Consider the following situation. Two people who are members of [the same culture or community/different cultures or communities] discuss whether **it is morally acceptable for a state to limit immigration.** One person says [—in line with the majority opinion of their culture—] that it is morally acceptable for a state to limit immigration. The other person says [—in line with the majority opinion of their culture—] that it is morally acceptable for a state to limit immigration. The other person says [—in line with the majority opinion of their culture—] that it is not morally acceptable for a state to limit immigration.

Which interpretation of this disagreement seems most appropriate to you?(1) One of these two people is correct and the other is incorrect.

- (2) Both people are correct.
- (3) Neither person is correct nor incorrect.

Responses were interpreted as follows: Intracultural (1) + Intercultural (1): Realism Intracultural (1) + Intercultural (2): Cultural relativism Intracultural (2) + Intercultural (2): Individual relativism Intracultural (3) + Intercultural (3): Non-cognitivism Others: Non-standard

Theory Task.

Is it morally acceptable for a state to limit immigration? Here, we are not interested in what is the correct answer to this question. Rather, we will present you with theories about whether there is a correct answer to the question, and if yes, what it is that makes this answer correct.

Which of these theories seems most appropriate to you?

- (1) There is a correct answer to the above question. This answer represents a fact. This fact is objective, i.e., independent from what anybody thinks about it. In other words, even if an individual or society were to regard the answer as incorrect it would still be correct. [REALISM]
- (2) There is a correct answer to the above question. This answer represents a fact. This fact is subjective, i.e., depends on what individuals or societies think about it. In other words, if an individual or society were to regard the answer as incorrect then this would make the answer incorrect. [RELATIVISM]
- (3) There is not a correct answer to the above question at all. Rather than representing a fact, any answer to this question just expresses a person's feelings, intentions, emotions or attitudes about it. [NON-COGNITIVISM]

Uncertainty Task.

Imagine that someone was uncertain whether **it is morally acceptable for a state to limit immigration**. We will present you with several interpretations of what it may mean to be uncertain about this question.

Which interpretation seems most appropriate/most obvious to you?

(1) The person does not know the objectively correct answer to this question.

[REALISM]

(2) The person does not know the answer that is considered correct by society.

[CULTURAL RELATIVISM]

- (3) The person does not know their own stance or opinion in this question. [INDIVIDUAL RELATIVISM]
- (4) The person has conflicting feelings, intentions, emotions or attitudes towards this question. [NON-COGNITIVISM]

Resolving Task.

Imagine that you yourself were uncertain whether it is morally acceptable for a state to

limit immigration.

How could you best reduce this uncertainty?

(1) I could look at what science or philosophy have to say about this question.

[REALISM]

(2) I could look at which answer to this question is given in my culture. [CULTURAL

RELATIVISM]

- (3) I could try to become clearer about my own opinion on this question. [INDIVIDUAL RELATIVISM]
- (4) This uncertainty could not be deliberately reduced. [NON-COGNITIVISM]

Study 2

Metaphor Task.

Is it morally acceptable for a state to limit immigration? Here, we are not interested in what is the correct answer to this question. Rather, we will present you with metaphors about whether there is a correct answer to the question, and if yes, what it is that makes this answer correct.

Which of these metaphors seems most appropriate to you?

- (1) There is a correct answer to this question. It is "out there in the world". [REALISM]
- (2) There is a correct answer to this question. It is "invented or created by society".

[CULTURAL RELATIVISM]

- (3) There is a correct answer to this question. It is "invented or created by each individual". [INDIVIDUAL RELATIVISM]
- (4) There is not a correct answer to this question. Any response to this question is just like someone saying "Booh!" or "Hurray!" [NON-COGNITIVISM]

Disagreement Task.

Disagreement task was the same as in Study 1.

Theory Task.

Is it morally acceptable for a state to limit immigration? Here, we are not interested in what is the correct answer to this question. Rather, we will present you with theories about whether there is a correct answer to the question, and if yes, what it is that makes this answer correct.

Which of these theories seems most appropriate to you?

(1) There is a correct answer to the above question. This answer represents a fact. This fact is objective, i.e., independent from what anybody thinks about it. In other words, even if an individual or society were to regard the answer as incorrect it would still be correct. [REALISM]
- (2) There is a correct answer to the above question. This answer represents a fact. This fact is subjective: It depends on what the respective culture or society thinks about it. In other words, if a culture were to regard the answer as incorrect then this would make the answer incorrect. [CULTURAL RELATIVISM]
- (3) There is a correct answer to the above question. This answer represents a fact. This fact is subjective: It depends on what the respective individual thinks about it. In other words, if an individual were to regard the answer as incorrect then this would make the answer incorrect. [INDIVIDUAL RELATIVISM]
- (4) There is not a correct answer to the above question at all. Rather than representing a fact, any answer to this question just expresses a person's feelings, intentions, emotions or attitudes about it. [NON-COGNITIVISM]

Uncertainty Task.

Imagine that someone was uncertain whether **it is morally acceptable for a state to limit immigration**. We will present you with several interpretations of what it may mean to be uncertain about this question.

Which interpretation seems most appropriate to you?

- Ultimately, the person is uncertain which answer to the question is objectively correct. [REALISM]
- (2) Ultimately, the person is uncertain which answer to the question is correct according to the standards of their own culture or society. [CULTURAL RELATIVISM]
- (3) Ultimately, the person is uncertain which answer to the question correctly reflects their own opinion or attitude. [INDIVIDUAL RELATIVISM]
- (4) The uncertainty does not refer to which answer to the question is correct. Instead, it is the result of conflicting feelings, intentions, emotions, or opinions about the question. [NON-COGNITIVISM]

Resolving Task.

Imagine that you yourself were uncertain whether **it is morally acceptable for a state to limit immigration**. Now you want to become more certain about this question and reduce your uncertainty. We will show you different descriptions of what you would need to do to reduce your uncertainty.

Which description seems most appropriate to you?

- (1) I would have to learn more about objective reality. [REALISM]
- (2) I would have to learn more about my culture or society. [CULTURAL RELATIVISM]
- (3) I would have to learn more about myself and my own views. [INDIVIDUAL

RELATIVISM]

(4) I would have to resolve my conflictedness on this question. [NON-COGNITIVISM]

Strategy Task.

Imagine that you yourself were uncertain whether **it is morally acceptable for a state to limit immigration**. Now you want to become more certain about this question and reduce your uncertainty. We will show you different strategies you could try to reduce your uncertainty.

Which strategy seems most appropriate to you?

- To listen to the opinion of experts (e.g. from science or philosophy) on this question.
 [REALISM NATURALISM]
- (2) To listen to my gut feeling or intuition about this question. [REALISM NON-NATURALISM]
- (3) To listen to opinions from society on this question. [CULTURAL RELATIVISM]
- (4) To reflect on my own views on this question. [INDIVIDUAL RELATIVISM]
- (5) The uncertainty could not be deliberately reduced. [NON-COGNITVISM]

A-3 Manuscript 3 – The Exculpatory Potential of Moral Ignorance: Evidence From a Blame-Updating Paradigm

Theisen, M. & Andow, J. (2024). *The exculpatory potential of moral ignorance: Evidence from a blame-updating paradigm*. Manuscript submitted for publication.

The Exculpatory Potential of Moral Ignorance: Evidence From a Blame-Updating

Paradigm

Maximilian Theisen¹, James Andow²

¹Department of Psychology, Heidelberg University, Heidelberg, Germany

²Department of Philosophy, University of Manchester, Manchester, UK

Author Note

This work was supported by an individual fellowship of the Baden-Württemberg Landesgraduiertenförderung Program and by a project grant from the heiDOCS-beCULT Doctoral Funding from Heidelberg University, Heidelberg, Germany. The authors have no competing interests to declare that are relevant to the content of this article. The data that support the findings reported in this article are openly available at

https://osf.io/9by25/?view_only=35369dde09d94a559aaa590e10c6275e.

Correspondence concerning this article should be addressed to Maximilian Theisen, Department of Psychology, Heidelberg University, Hauptstraße 47-51, 69120 Heidelberg, Germany. Email: maximilian.theisen@psychologie.uni-heidelberg.de

Abstract

Over the past few decades, philosophers have started to scrutinize the epistemic conditions on moral responsibility. One central focus of this debate is whether ignorance of moral norms can excuse wrongdoing in the same way as ignorance of the factual circumstances of an action can excuse wrongdoing. Volitionists link the exculpatory potential of moral ignorance to the fulfillment of the agent's procedural obligations. Quality of Will theorists have suggested the further conditions of the inaccessibility of the moral truth and high moral difficulty. In a preregistered experiment (*N* = 251), we tested whether these three conditions reduce ascriptions of blame for everyday moral transgressions (act blameworthiness) and for moral ignorance about their wrongness (belief blameworthiness). Utilizing linear mixed models, path analysis, and latent change score modeling, we did not find any evidence of a substantial exculpatory effect of moral ignorance. Even though the data suggest that moral ignorance might be perceived as exculpatory if it were fully non-culpable, we did not observe such an effect as participants did not judge any of our scenarios to involve fully non-culpable moral ignorance. We discuss the relevance of our findings for debates about epistemic conditions on moral responsibility and for discussions of choice under moral uncertainty.

Keywords: blame; moral responsibility; epistemic condition; moral ignorance; normative externalism; normative internalism

Introduction

Philosophers widely agree that being blamelessly ignorant of the wrong-making features of an action can excuse wrongdoing (Robichaud & Wieland, 2017) – a view also identified in the lay public (Kirfel & Hannikainen, 2023; Kissinger-Knox et al., 2018; Monroe & Malle, 2019). For ignorance about moral norms themselves, however, the verdict is less clear. Some argue that moral ignorance can exculpate in the same way and under the same conditions as factual ignorance. Others deny this, either because they see moral ignorance in itself as sufficient for blame, or because they doubt that the relevant conditions can be met in cases of moral ignorance. We tested whether the conditions for exculpatory moral ignorance that have been discussed in the literature would be sufficient for people to reduce ascriptions of blame for morally ignorant agents.¹

Exculpatory Ignorance

Boris keeps a certain species of exotic animals at his home. Unbeknownst to him, the conditions in which he holds these animals cause them intense distress, leading to specific behavior in the animals. Boris mistakenly interprets this behavior as a sign of excitement and therefore thinks that he is treating the animals right. Is Boris to blame for keeping the animals in distressing conditions?

The answer that is often given is that the agent's blameworthiness depends on the culpability of the ignorance itself (Robichaud & Wieland, 2017; but see Hartford, 2019; Smith, 1983). We can easily imagine further details about Boris's case that would increase or reduce his blameworthiness depending on the circumstances of his ignorance. Supposing that Boris has

¹ In philosophy, excuses are often understood as ridding the agent of any responsibility. We will follow this usage in our discussion of the philosophical literature. However, from an empirical perspective, we will understand excuses as gradual reductions in blameworthiness (Sliwa, 2020).

never attempted to learn about the behavior and habits of the species he keeps, it seems clear that he is to blame for keeping the animals in such conditions. However, supposing that Boris's interpretation of the animals' behavior as excitement is based on what was, until very recently, the best science of animal behavior and that the latest findings are yet to be effectively communicated to people in Boris's position, we might be inclined to let him off the hook.

So far, we assumed that Boris's ignorance was purely factual. Let's now turn to a case where the agent is morally ignorant. Taylor is in a very similar situation as Boris. He holds the same species of animals at his home and the conditions he offers them also cause the animals to be immensely distressed. In contrast to Boris, Taylor is aware of this. However, Taylor does not think that it is wrong of him to hold the animals in such conditions. Even though Taylor is aware of the wrong-making features of his action, namely the suffering that he causes the animals, he is not aware *that these features are wrong-making*. Is Taylor to blame for keeping the animals in distressing conditions?

Taylor's case raises the question of whether moral ignorance holds exculpatory potential, just as factual ignorance does. In the following, we will briefly sketch two contrasting views about this question.

Volitionism and Procedural Obligations

The current debate regarding exculpatory moral ignorance largely goes back to volitionist arguments as they have been laid out in detail by Rosen (2003, 2004) and Zimmerman (1997). Rosen argues that people have obligations regarding the management of their factual and moral beliefs with the aim of avoiding the risk of negligent wrongdoing (*procedural obligations*). In his view, an agent can *only* ever be blameworthy for ignorant wrongdoing if the agent was aware of their procedural obligations but failed to act accordingly (*akrasia*), leading to blameworthy ignorance. For Boris, these procedural obligations might have involved informing himself on the conditions his animals require. For Taylor, these obligations might have involved

reflecting on the interests of the animals or engaging in discussion with his peers who he knows to hold different moral views on animal rights than him.

It is usually assumed that for many cases of ignorance, the criterion of akrasia is not met (e.g., FitzPatrick, 2008; Rosen, 2003, 2004). Especially in cases of moral ignorance, the ignorance might instead be the result of the agent's upbringing or general social environment, without them ever becoming aware of their obligation to question their moral views (Rosen, 2003; Wolf, 1987). Volitionists interpret this as suggesting that most people who act out of ignorance are blameless for their wrongdoing. Non-volitionists, on the other hand, take this as indicative of akrasia being too strict a criterion for moral responsibility.

Quality of Will, Moral Evidence, and Moral Difficulty

Many of those who argue that Rosen's view is too lenient, as it lets apparently blameworthy agents off the hook, do so from a *Quality of Will* (QW) perspective (Arpaly, 2003; Harman, 2011, 2022; see also Hartford, 2019). Proponents of QW argue that an agent is blameworthy for their wrongdoing if their action arises from bad will, such as inappropriate concern for the morally-relevant aspects of that action. For example, Boris would be considered blameworthy if his ignorance and consequently his wrongdoing would be a manifestation of his lack of concern for the animals.

For ignorance to be exculpatory in this view, it would have to disrupt the link between the wrongdoing and the agent's moral concerns. In our example, if Boris's treatment of his animals was based on what he had to assume was the current state of science, it would not arise from a lack of concern for the animals and he would be excused for it. However, if Boris has never attempted to inform himself on what his animals require, his ignorance could be understood as illuminating his lack of concern for those animals, and he would therefore not be excused by it.

The case of Boris shows that volitionism and QW can lead to similar conclusions regarding the blameworthiness of factually ignorant agents. In the case of moral ignorance,

however, both views can come apart. If we consider the case of Taylor, volitionism would suggest that Taylor is blameless for his treatment of the animals, as long as his moral ignorance was not the result of his wittingly violating his procedural obligations. From a QW perspective, however, Taylor would typically be considered blameworthy, even if he was not responsible for his ignorance. Taylor's not being aware of the relevance of the animals' suffering could in itself be seen as a manifestation of his lack of appropriate concern, and therefore suffice to hold him accountable.

Based on such considerations, some proponents of QW have argued that moral ignorance can never be exculpatory. This view has been defended in detail by Harman (2011, 2022). Harman's argument rests on the assumption that moral ignorance is in itself always a manifestation of bad will, thereby making it in principle unsuitable for exculpation. Others have argued that there might be conditions where moral ignorance is not revelatory of the agent's quality of will, leaving room for moral ignorance to sometimes be exculpatory (Hartford, 2019; Wieland, 2017). These arguments in favor of the compatibility of QW and exculpatory moral ignorance have focused on the accessibility of the respective moral information. If the agent has fulfilled their procedural obligations but the moral truth has not been accessible to them, the agent's ignorance is outside of their control and they would not be blameworthy according to this reading of QW. Two ways that a well-meaning agent might lack access to the moral truth have been discussed in the literature.

First, the moral truth might not have been accessible because the moral evidence that was available to the agent, such as testimony or first-hand experience, was biased towards a false moral view (Arpaly, 2003; Hartford, 2019; Wieland, 2017). For example, it might be that Taylor grew up in an environment where animals were mostly treated as commodities rather than fellow sentient beings. Further, he might be surrounded by people who share his views on the issue, not to mention all his friends who also hold animals in suboptimal conditions without anyone expressing disapproval towards them. In such a case, it would have been particularly hard for Taylor to recognize that what he is doing is wrong and it might be that his ignorance would be exculpatory due to that.

Another way in which moral truth might not be accessible is by high difficulty of the moral issue itself, such as high complexity of the morally relevant aspects, a high level of conflict between competing values, etc. (Bradford, 2017; FitzPatrick, 2008; Guerrero, 2017; Harman, 2022; Hartford, 2022). While a false belief regarding easy moral cases might always be revelatory of the agent's will, so the idea, maybe moral ignorance about complex ethical issues is not, or at least not to the same degree. First, reasoning about difficult moral cases can be "indeterministic" (Rosen, 2003, p. 70), so that coming to the right conclusions is not truly in the hand of the agent, even if they fulfill their procedural obligations. Second, cases of high moral difficulty are often characterized by multiple, conflicting moral values that are at stake (Makins, 2021). An agent might be morally ignorant about such cases by giving too much weight to one of these values, thereby partially neglecting the other. Such ignorance, one might argue, is less revelatory of the agent's lack of concern for what is morally relevant than if the agent were ignorant of one of these values. The latter would reveal their absolute disregard of the respective value. The former, however, would merely reveal that they underestimate the relevance of a specific aspect *relative to another* (see also Biebel, 2023).

In conclusion, QW is typically understood to rule out exculpatory moral ignorance, but some have attempted to reconcile both ideas by focusing on cases where the agents are "more ignorant than [morally] vicious" (Arpaly, 2003, p. 304; Hartford, 2019, p. 1101), trying to undermine the link between moral ignorance and bad will. The two potential conditions of exculpatory moral ignorance that have emerged from the discussion of QW are a lack of suitable moral evidence and high moral difficulty.

Currently, the debate has come to a stalemate between those that see moral ignorance as exculpatory and those that do not. We hope to enrich the debate on the basis of empirical findings regarding the intuitiveness of the arguments that so far have been provided in favor of or against exculpatory moral ignorance.

Previous Findings Regarding Exculpatory Moral Ignorance

Previous studies clearly find that the exculpatory potential of factual ignorance is not only recognized by most philosophers but also by lay people (Kirfel & Hannikainen, 2023; Kissinger-Knox et al., 2018; Monroe & Malle, 2019). Regarding moral ignorance, the evidence is sparser. In one of the first investigations into this topic, Faraci and Shoemaker (2010) used variations of the seminal JoJo case (Wolf, 1987), describing the morally compromising upbringing of a dictator's son, to test whether moral ignorance resulting from such an upbringing would lead to lower ascriptions of blame than moral ignorance by itself. The idea of cases like that of JoJo is that a morally compromising upbringing describes circumstances where the moral truth is not accessible to the agent. Faraci and Shoemaker found that moral ignorance resulting from morally compromising upbringing led to lower ascriptions of blame than moral ignorance by itself (moderate effect size). Interestingly, whether the dictator's son had been exposed to alternative moral views did not affect his perceived blameworthiness. Faraci and Shoemaker (2014) replicated their findings utilizing the case of a contemporary racist who was either described as forming his false moral view during his adulthood or as the result of his upbringing. This second study found a large exculpatory effect of morally compromising upbringing. Both studies have been interpreted as showing the relevance of the accessibility of moral truth for the exculpatory potential of moral ignorance. However, it should be noted that both study designs did not allow to differentiate between effects of moral ignorance resulting from morally compromising upbringing and effects of the upbringing itself (we elaborate on this aspect in the discussion of our own findings).

More recently, Schwartz and Inbar (2023) investigated how descriptions of the moral beliefs of an agent affected ratings of act wrongness and moral character in the context of everyday moral transgressions. They found that actions resulting from moral ignorance were

rated to be just as wrong as the control cases, but cases of akrasia, where the agent knew that their action was impermissible, were rated as more wrong than the others. Regarding moral character, agents showing akrasia were rated more positively than the control and morally ignorant agents were rated more negatively than the control. These findings suggest that lay people evaluate acts of akrasia negatively and that moral ignorance by itself is not perceived as reducing or elevating act wrongness. Further, they show that acts of akrasia reflect more positively on the moral character of the agent than acts committed under moral ignorance, suggesting that moral character ratings are closely linked to ratings of the moral beliefs of the agent.

Moral responsibility was not the main focus of Schwartz and Inbar's work. Consequently, they did not consider the conditions that have been discussed as necessary for exculpatory moral ignorance nor were they interested in ascriptions of blameworthiness. In the current study, we addressed the shortcomings of previous studies by explicitly testing the exculpatory effect of different variations of moral ignorance based on arguments from the philosophical literature. We investigated their effect on act blameworthiness (i.e., how blameworthy the agent is for their wrongdoing) and belief blameworthiness (i.e., how blameworthy the agent is for their ignorance) to gain a full picture of the exculpatory potential of moral ignorance and its underlying dynamics in the lay public.

Experiment

We tested whether philosophical lay people would update their blame judgments after learning that the agent was morally ignorant. Taking inspiration from Schwartz and Inbar (2023), we focused on blame for everyday moral transgressions. This stands in contrast to most of the thought experiments in the philosophical literature which feature idiosyncratic circumstances to boost intuitions in favor of exculpation, such as cases of ancient human trafficking, mid-century sexism, or the previously mentioned case of JoJo (Rosen, 2003; Wolf, 1987). These cases are designed to show how blameless the respective agent is for their ignorance by putting them in a social environment that supports their false moral view. However, as we previously mentioned, such cases might not be suited to differentiate between effects of moral ignorance and effects of the environment itself. We therefore decided to present cases of agents who might be part of everyday life.

We assumed that any exculpatory effect of moral ignorance would be closely linked to how blameless the agent is for being ignorant. We focused on the three conditions of exculpatory moral ignorance that we have discussed above: First, following volitionism, an agent might be blamelessly ignorant if they have shown due care in the management of their moral beliefs by fulfilling their procedural obligations (Rosen, 2003, 2004). This means that they have done everything that can be reasonably expected from them to get the moral evaluation of the respective issue right, such as considering the relevant available arguments or questioning their own sentiments and intuitions regarding the issue. In this view, the moral mistake would happen somewhere in the considerations of the agent, but outside of their immediate control. We call this the *fulfilled procedural obligations* condition.

Second, following a moderate version of QW, an agent might be blameless for having a false moral view if the view was supported by the available moral evidence, including testimony and first-hand experiences, so that the view was not the result of flawed moral reasoning at all (Hartford, 2019; Wieland, 2017). We call this the *biased moral evidence* condition.

Finally, we tested whether the exculpatory potential of moral ignorance would depend on the moral difficulty of the respective issue (Bradford, 2017; FitzPatrick, 2008; Guerrero, 2017; Harman, 2022; Hartford, 2022). We assumed that having a false moral belief about a difficult moral case would be perceived as less culpable than having a false moral belief about an easy moral case. We call this the *moral difficulty* condition.

We investigated the effects of these conditions on the exculpatory power of moral ignorance using a blame-updating paradigm (Monroe & Malle, 2019). In this paradigm, participants are confronted with the description of an action and are asked for an initial blame

judgment. Then, they receive further information and get the opportunity to change their initial response. Using such a paradigm, Monroe and Malle could show that lay people rationally update their blame judgments in the light of new information, such as learning that the agent acted unintentionally. The exculpatory effects identified by Monroe and Malle served as a benchmark for the effects in our study.

Hypotheses

Hypotheses and analysis strategy were preregistered on OSF

(https://osf.io/97yse/?view_only=fb44ab1b66a849599ba67323352f4f64). We started with an investigation of differences between control trials and moral ignorance trials. We expected that moral ignorance would lead to a reduction of act blameworthiness in comparison to the control trials (*H1*). We also expected that this effect would be more pronounced in moral issues that were perceived as difficult in comparison to easy moral cases (*H2*). Further, we expected that an overall exculpatory effect of moral ignorance would be solely the result of cases where both fulfilled procedural obligations and biased moral evidence were given (*H3*). This was tested by comparing each moral ignorance condition against the control condition.

Focusing on moral ignorance trials only, we expected that fulfilled procedural obligations, biased moral evidence, and moral difficulty would predict changes in act blameworthiness (*H4*). Following both volitionist and QW arguments, we further assumed that changes in act blameworthiness would be closely associated with belief blameworthiness, i.e., as how blameworthy the agent was perceived for being ignorant (*H5*), and that the effects identified for *H4* would be mediated by belief blameworthiness (*H6*).

Method

Participants

We recruited 251 participants (161 male, 81 female, 4 non-binary, 5 not indicating) in the age between 18 and 73 (M = 30.86, SD = 9.72) via *Prolific*. Participants were required to be fluent in German, have an approval rate of at least 95%, and be located in Germany, Switzerland, or Austria. Many participants had a bachelor's degree or higher (N = 116), or had an otherwise high educational degree (N = 102). Participants were paid £2.5 for their participation which was expected to take approximately 14 minutes. Actual median completion time was 9:56 mins.

Design and Procedure

Stimulus Categorization and Moral Difficulty. The study consisted of two phases: (1) The stimulus selection including ratings of moral difficulty, and (2) the blame-updating task including ratings of act wrongness, act blameworthiness, and belief blameworthiness. For stimulus selection, participants were shown a list of 39 actions in randomized order and were asked to select five actions that they morally tended to reject or find wrong (Table 1). On a separate page, participants were asked to confirm their selection or to correct it where necessary. Participants were then asked to indicate how certain they are in their moral evaluation of each of the selected actions on a scale from 0 (*not certain at all*) to 100 (*absolutely certain*). We calculated a proxy for perceived moral difficulty of the respective action as 100 – certainty rating.

Blame-Updating Task. After stimulus categorization, participants started the blameupdating task where they were asked to rate persons that often commit one of the selected actions. There were five trials for each participant: Four moral ignorance trials following a 2 (procedural obligations: fulfilled vs. violated) × 2 (moral evidence: biased vs. unbiased) design and one additional control trial. The order of trials was randomized. The five actions that had been selected in the stimulus categorization task were randomly assigned to these trials.

Each trial of the blame-updating task consisted of two screens. On the first screen, participants received basic information about the agent's name and the action (e.g., "Nick often eats meat"). They were asked to rate act blameworthiness and act wrongness. All measurements used a slider scale as response format. Act blameworthiness was assessed using three items: "How much is [the agent] morally responsible for this behavior?" (0 = not at all responsible, 100 = fully responsible); "How much is [the agent] to blame for this behavior?" (0 = *no blame at all*, 100 = *maximum blame*); "How much reproach does [the agent] deserve for this behavior?" (0 = *no reproach*, 100 = *maximum reproach*). Act wrongness was assessed on a single item: "In general, how immoral do you think it is to [do the action]?" (0 = *not immoral at all*, 100 = *completely immoral*).

On the second screen, participants received further information about the agent, depending on the experimental condition. In the control condition, they read: "Nick often eats meat. It is not known whether he thinks this action to be morally acceptable or inacceptable". The control trial was used to control for retest effects and for effects of merely mentioning that the agent might have a moral belief concerning the respective moral issue. In the moral ignorance trials, the agent was characterized as thinking of their action as morally permissible, which was further characterized according to the experimental manipulation of procedural obligations and available moral evidence (Table 2). After reading the new information about the agent, participants were asked again to rate act blameworthiness and act wrongness. They were told that they could change or retain their initial judgment. Participants could not revisit previous sites and their previous ratings of act blameworthiness and act wrongness were not shown to them on the second screen of the blame-updating task. In the moral ignorance trials, but not in the control trials, participants were further asked to rate the agent's belief blameworthiness: "How much is [name] responsible for having arrived at this moral evaluation of the issue?" (0 = not at all responsible, 100 = fully responsible); "How much is [name] to blame for their moral view?" (0 = no blame at all, 100 = maximum blame); "How much reproach does [name] deserve for their moral view on the subject?" (0 = no reproach, 100 = maximum reproach).

Study End. After the blame-updating task, participants responded to questions about their gender, age, and education. They were further asked whether they had responded conscientiously. They were informed that their response to this question would not have any consequences for them. Finally, participants were given the opportunity to comment on the

study, received information about the aim of the study, and were given their completion code for compensation.

Analysis

Hypotheses 1 - 5 were tested in linear mixed models (for an overview of the fitted models, see Table 3). Hypothesis 6 was tested using path analysis and structural equation modeling.²

We built composite scores for act blameworthiness and belief blameworthiness. We calculated change scores for act blameworthiness and act wrongness by subtracting the prescores from the respective post-scores.³ As a robustness check, we confirmed the findings from the manifest path analysis using latent change score modeling in a structural equation modeling context (Kievit et al., 2018; Steyer et al., 1997). The theoretically assumed model is described in

² Linear mixed models were fitted using the *Ime4*-package (version 1.1.31; Bates et al., 2015) in *R* (version 4.2.2; R Core Team, 2022). For path analysis and structural equation modeling, we used the *lavaan*-package (version 0.6.14; Rosseel, 2012). To account for the hierarchical nature of our data, *lavaan* models were fitted with cluster-robust standard errors with participants as clusters (McNeish et al., 2017). As estimator, we chose MLR which is suitable for non-normal data. Structural equation models were estimated based on standardized variables.

³ There is an ongoing debate on the appropriateness of using manifest change scores as dependent variables and how this procedure compares to competitive approaches, such as ANCOVA or residual change scores. Simulation studies suggest that in an experimental setting with full randomization, as it is the case in our study, type 1 error rates, power, and bias of change scores are comparable to those of their competitors (Jennings & Cribbie, 2016). Additionally, change scores exhibit high interpretability in the context of blame updating (Monroe & Malle, 2019). Therefore, we chose manifest change scores for our analyses. Figure 1. *Lavaan* syntax of the final model can be found in the Supplementary Material of this article (Supplementary Code S1).

Results

Measurement Characteristics

The data underlying our analyses are available at OSF

(https://osf.io/9by25/?view_only=35369dde09d94a559aaa590e10c6275e). Overall, participants gave high ratings of act and belief blameworthiness, and low ratings of moral difficulty (Table 4). Reliability estimates for the composite scores for act ($\omega_{between} = .902$, $\omega_{within} = .782$) and belief blameworthiness ($\omega_{between} = .941$, $\omega_{within} = .816$) were good.

Moral Ignorance vs. Control Trials

Were There Any Changes in Ratings of Act Wrongness? There was a possibility that the retesting itself or reading about an agent who believes their action to be acceptable might influence participants' evaluations of act wrongness (Schwartz & Inbar, 2023). In the case of such effects, we would have needed to control for them when testing for effects on blameworthiness. We therefore started by testing whether there were any pre-post changes in ratings of act wrongness using a linear mixed model (Model 1a). The inclusion of a random intercept for action led to singularity issues, so we only included a random intercept for participant (Model 1b). We did not find any significant changes in act wrongness in the control trials, b = -0.19, 95% CI [-1.21, 0.83], t(1211.843) = -0.369, p = .713, d = -0.023, 95% CI [-0.147, 0.101], nor was there any main effect of moral ignorance trials, b = 0.39, 95% CI [-0.69, 1.48], t(1003) = 0.713, p = .476, d = 0.048, 95% CI [-0.014, 0.11]. Using Model 1c, we also did not find any changes in act wrongness ratings in the moral ignorance trials, b = 0.2, 95% CI [-0.37, 0.78], t(337.733) = 0.694, p = .488, d = 0.025, 95% CI [-0.037, 0.087]. In conclusion, act wrongness ratings were unaffected by the presentation of information about the agent's moral ignorance and by the repeated measurement itself. They were thus not considered in the further analyses.

Did Moral Ignorance Have an Overall Exculpatory Effect? To test for an overall exculpatory effect of moral ignorance, we tested whether the difference between pre- and postmeasurements of act blameworthiness differed between moral ignorance and control trials. We fitted a linear mixed model predicting changes in act blameworthiness from moral ignorance versus control trials, including crossed random intercepts for participant and action (Model 2a). The intercept for predicting changes in act blameworthiness was significantly negative, indicating that merely repeating the measurement and mentioning that the moral beliefs of the agent are not known (control condition) slightly reduced ascriptions of blame from pre- to postmeasurements, b = -1.51, 95% CI [-2.85, -0.17], t(207.714) = -2.214, p = .028, d = -0.147, 95% CI [-0.271, -0.023]. The main effect of moral ignorance was positive and significant, albeit of similarly small magnitude, b = 1.43, 95% CI [0.09, 2.76], t(998.112) = 2.097, p = .036, d = 0.148, 95% CI [0.005, 0.292]. Model 2b revealed that both of these effects cancelled each other out, so that there wasn't any significant blame updating in the moral ignorance trials, b = -0.08, 95% CI [-0.96, 0.77], *t*(35.332) = -0.179, *p* = .859, *d* = -0.014, 95% CI [-0.076, 0.047]. An overall exculpatory effect of moral ignorance could therefore not be stated and H1 was not confirmed. Adding moral difficulty (Model 2c) or moral difficulty and its interaction with moral ignorance (Moral 2d) to the model did not improve model fit any further. Change in act blameworthiness was not affected by moral difficulty, neither in general nor specifically in the control or moral ignorance trials. H2 was therefore also not confirmed.

Did Fulfilled Procedural Obligations and Biased Moral Evidence Lead to an Exculpatory Effect of Moral Ignorance? We could rule out that there was an overall exculpatory effect of moral ignorance. However, we had hypothesized that moral ignorance would only be exculpatory under specific conditions, i.e., when procedural obligations were fulfilled and the available moral evidence was biased towards the false moral view. Such an effect was not ruled out by the absence of an overall exculpatory effect. We therefore continued by testing each combination of the procedural obligations and the moral evidence factors against the control condition (Model 3). To control for the multiple testing, we used Dunnett contrasts (Figure 2).

The only blame reduction that was significantly different from zero was found for fulfilled procedural obligations in combination with biased moral evidence, which was in line with our theoretical reasoning. However, this reduction did not differ significantly from that in the control condition. The exculpatory effect of this factor combination could therefore not be differentiated from a mere retest effect and *H3* was not confirmed.

Instead, there was a small but significant increase in blame ascriptions in those conditions with unbiased moral evidence in comparison to the control condition. This indicated that agents were blamed slightly more for their actions if they had been exposed to the right moral view before acquiring a false moral view of the respective issue.

In conclusion, the comparisons between moral ignorance and control trials revealed that moral ignorance was neither seen as generally exculpatory nor as exculpatory under the different combinations of procedural obligations and moral evidence that we had tested in this study. The effects that could be identified in these comparisons were all of small magnitude. As Figure 2 shows, the largest change between pre- and post-measurements of act blameworthiness could be seen in the combination of fulfilled procedural obligations and biased moral evidence. However, this difference was more than four times smaller than the exculpatory effect that had been identified by Monroe and Malle (2019) for when the wrongdoer was described as acting unintentionally, and it did not significantly differ from the control condition.

Effects Within Moral Ignorance Trials

Did Fulfilled Procedural Obligations, Biased Moral Evidence, and Moral Difficulty Affect Changes in Act Blameworthiness? Even though none of the experimental cells had revealed an exculpatory effect that differed from the control, the results so far suggested differences in blame updating between the different conditions of moral ignorance. To get a fuller picture of main and interaction effects, we fitted a linear mixed model with procedural obligations, moral evidence, their interaction, and moral difficulty as predictors of change in act blameworthiness (Model 4a). There was a significant main effect for moral evidence and a significant interaction between moral evidence and procedural obligations, but no main effect for procedural obligations or moral difficulty. We therefore reran the model without moral difficulty, which improved model fit (Model 4b). The resulting model suggested that within the moral ignorance trials, blame was reduced if the agent had reasoned under biased moral evidence, with an additional blame reduction if the biased moral evidence was combined with a fulfillment of the agent's procedural obligations (Table 5). *H4* could therefore be partially confirmed; however, the magnitude of the identified effects was again small.

Were Changes in Act Blameworthiness Associated With Belief Blameworthiness? It has often been argued that an exculpatory effect of moral ignorance would require that the ignorance itself is non-culpable, suggesting a positive association between belief blameworthiness and changes in act blameworthiness. This association was confirmed in our data, so that participants were more likely to reduce ascriptions of blame if they found the agent less blameworthy for their moral ignorance, b = 0.23, 95% CI [0.19, 0.27], t(759.059) = 11.03, p <.001, $R^2_{partial} = 18.7\%$ (Model 5). The effect size was moderate to large. *H5* was confirmed.

The association between both variables suggested that a considerable blame reduction could be expected if the agent were fully blameless for their ignorance (Figure 3). However, in the vast majority of cases, the agent was held responsible for their ignorance in our study.

Did Fulfilled Procedural Obligations, Biased Moral Evidence, and Moral Difficulty Affect Belief Blameworthiness? Having established that belief blameworthiness was highly relevant for an exculpatory effect of moral ignorance, we then tested how procedural obligations, moral evidence, their interaction, and moral difficulty in turn affected belief blameworthiness (Model 6a). In this model, only moral evidence and moral difficulty were significant. However, model fit was not improved by removing the non-significant predictors (Model 6b & Model 6c). We therefore decided to interpret the full Model 6a (Table 6). Participants blamed the agent less for their moral ignorance if the participants perceived the moral issue to be difficult (moderate effect size) and, to a lesser degree, if the agent had formed their ignorance as a result of biased moral evidence (small effect size) (Figure 4). Note that the investigated conditions were not sufficient for participants to consistently see the agent as blamelessly ignorant. Even when assuming maximally difficult moral cases, which were rare in our study, the estimated belief blameworthiness was still at around 60 out of 100 points (Figure 4B).

Were Effects on Changes in Act Blameworthiness Mediated by Belief

Blameworthiness? We assumed that fulfilled procedural obligations, biased moral evidence, their interaction, and high moral difficulty would affect changes in act blameworthiness by affecting belief blameworthiness. One necessary assumption of such a causal structure, namely a respective correlational structure, was tested utilizing mediation analysis.⁴ We tested the assumed mediations both on a manifest level using path analysis, as well as on a latent level using latent change score modeling in a structural equation modeling context.

Path Analysis. Just as in the linear mixed models, path analysis did not reveal any effect of fulfilled procedural obligations, neither on belief blameworthiness nor on changes in act blameworthiness (Figure 5 & Table 7). We found that the effect of biased moral evidence was fully mediated by belief blameworthiness, as indicated by a significant indirect effect and a nonsignificant direct effect. The interaction of fulfilled procedural obligations and biased moral evidence, on the other hand, was not mediated by belief blameworthiness, indicating that fulfilled procedural obligations in combination with biased moral evidence affected the

⁴ Note that mediation analysis is not a sufficient test of causality since the identified correlational structure might result from different causal structures. Causality needs to be derived from study design and theoretical arguments (Fiedler et al., 2011).

exculpatory potential of moral ignorance without affecting how blameworthy the agent was perceived for their ignorance.

The previous analyses had suggested that moral difficulty did not have any effect on changes in act blameworthiness. However, path analysis revealed a suppressor effect: There was an indirect, blame-mitigating effect of moral difficulty that was mediated by belief blameworthiness. However, there was also a direct blame-enhancing effect of moral difficulty, that was not in line with our theoretical assumptions. The two conflicting effects cancelled each other out, leading to a non-significant total effect of moral difficulty, which explains its null effect on changes in act blameworthiness in the previous analyses. Overall, *H6* was partially confirmed, with a fully mediated effect of biased moral evidence, and a suppressed indirect effect of moral difficulty. Against our expectations, the interaction effect of fulfilled procedural obligations and biased moral evidence was not mediated by belief blameworthiness.

Latent Change Score Modeling. To test the robustness of our findings from path analysis, we used structural equation modeling which allows controlling for measurement error and modeling true change between pre- and post-measurements on a latent level, resolving potential reliability problems of manifest change scores (Kievit et al., 2018; Steyer et al., 1997). First, we tested whether pre- and post-measurements of act blameworthiness measured the same latent construct by progressively fixing factor loadings, intercepts, and measurement errors to be equal between both points of measurement. We found evidence of strong measurement invariance, i.e., equal factor loadings and intercepts but unequal measurement errors, which allowed us to continue with the latent change score modeling (Supplementary Table S1).

A full model (Figure 1) including all predictors, belief blameworthiness as mediator, and a latent change score parameter showed good fit to the data, $CFI_{robust} = 0.978$, $TLI_{robust} = 0.969$, $RMSEA_{robust} = .056$, 90% CI [.046, .066], SRMR = 0.041. Fixing insignificant paths to zero did not improve the model, $CFI_{robust} = 0.976$, $TLI_{robust} = 0.971$, $RMSEA_{robust} = .054$, 90% CI [.045, .064], *SRMR* = 0.044. However, model fit could be improved by fixing insignificant intercepts to zero, $CFI_{robust} = 0.979$, $TLI_{robust} = 0.973$, *RMSEA*_{robust} = .052, 90% CI [.042, .062], *SRMR* = 0.041. The interpretation of results did not differ substantially between these model specifications. We decided to report the latter model due to its superior fit to the data. Full parameter estimates for this model can be found in the Supplementary Material of this article (Supplementary Table S2). Overall, the direct, indirect, and total effects identified in the manifest model were all confirmed on the latent level, except for the total effect of biased moral evidence which closely missed significance in the latent change score model (p = .050) (Figure 6 & Table 7).

Free Form Comments

The general theme of our findings was mirrored in free form comments from the study participants. These are the comments that directly referred to our hypotheses (translated from German by one of the authors):

- "I think that you are responsible for your moral views."
- "I find people who have engaged with the issue are of course fully responsible for their behavior. Likewise, I find people who have hardly engaged with it must also bear full responsibility. It would have been their duty to concern themselves with it a little more."
- "I didn't care what kind of reasoning was behind their actions. Certain actions [...] are simply morally reprehensible."
- "I have often wondered which is worse: when you strongly question your position and then still [commit the wrongdoing], or, when you have not thought about your actions. I don't know."

Discussion

Overall, moral ignorance did not lead to a reduction of blame ascriptions in our study. Even though there were small effects of biased moral evidence, its interaction with fulfilled procedural obligations, and high moral difficulty, those were not sufficient to create an exculpatory effect. There was a strong association between belief blameworthiness and changes in act blameworthiness, suggesting that moral ignorance might have substantial exculpatory power if the agent was fully blameless for their ignorance. Yet, this burden proved high since neither the experimental conditions nor high moral difficulty were sufficient to seriously undermine the sense that the agent was accountable for their false moral view. Overall, our findings are in line with those that accept that blameless moral ignorance might exculpate in principle but who doubt that conditions for blameless moral ignorance are easily met in reality.

Findings Regarding Procedural Obligations

Our findings put pressure on the intuitiveness of the volitionist argument that a necessary condition for the blameworthiness of an ignorant wrongdoer is the witting violation of procedural obligations. The findings of the current study strongly suggest that this view is not in line with how lay people ascribe blame. Having fulfilled or violated one's procedural obligations did, by itself, not play any role for blame-updating in our study.

It should be acknowledged that Rosen (2003) does not argue that volitionism is in line with common sense or everyday moral practice. However, he argues for the intuitiveness of volitionism once the respective arguments and examples have been put forward. In our study, even though participants were not confronted with volitionist arguments directly, they were made explicitly aware of the agent's thought process, how the agent managed their moral beliefs, and the resulting moral ignorance. Since this did not affect the amount of blame participants ascribed to the agent whatsoever, the current data did not lend any support to a volitionist perspective on moral ignorance. As we see it, the ball is now in the volitionists' court.

Findings Regarding Moral Evidence and Moral Difficulty

Proponents of QW have argued that fulfilled procedural obligations might not be sufficient to get a morally ignorant wrongdoer off the hook. Instead, it has been proposed that fulfilled procedural obligations need to be combined with low accessibility of the moral truth. In our study, we focused on two aspects of moral accessibility: the moral evidence that was available to the agent and the overall moral difficulty of the respective issue as perceived by the individual participant.

Moral Evidence

We could identify a blame-reducing effect of biased moral evidence, i.e., of whether the agent had only encountered arguments and experiences that supported their false moral view. This effect was mediated by belief blameworthiness and was of small magnitude. A slightly stronger effect could be identified for the interaction of fulfilled procedural obligations and biased moral evidence. In fact, this was the only experimental cell where the pre-post difference was significantly negative, indicating a small exculpatory effect which, however, was not different from that in the control trials. Interestingly, this interaction was not mediated by belief blameworthiness, which suggested that the associated blame reduction was not due to the agent being perceived as less blameworthy for their moral ignorance.

The lack of a substantial exculpatory effect in the biased moral evidence condition conflicts with the findings from Faraci and Shoemaker (2014) who had found a large exculpatory effect for moral ignorance resulting from morally compromising upbringing. A morally compromising upbringing is typically understood as a prime example of biased moral evidence. It might be that our general and abstract descriptions of the available moral evidence as biased and one-sided were not sufficient to elicit the respective intuitions in our participants. Perhaps, lay people see moral truth as easily accessible, similar to Harman (2022), so that it might have required concrete descriptions of how the moral truth was hidden from the agent to undermine this intuition. Following this interpretation, there is an exculpatory effect of moral ignorance resulting from biased moral evidence, as evident by the findings of Faraci and Shoemaker, but the threshold for counting as sufficiently biased was not reached in our study.

Another explanation for the strong effects identified by Faraci and Shoemaker might be that they described their ignorant agent as thinking that their wrongdoing is morally required. In our study, the agent was simply described as thinking their wrongdoing to be acceptable (see also Wieland, 2017).

As previously discussed, it might also be that the effects identified by Faraci and Shoemaker are effects of the agent's social environment itself, not of the resulting moral ignorance, which would also explain why we did not find effects of similar magnitude in our study. We see two ways in which social environment might have a direct effect on blame ascriptions independent of moral ignorance. First, descriptions of morally compromising environments might fundamentally shift our perception of the agent. Growing up in an environment where false moral values, such as racist beliefs (Faraci & Shoemaker, 2014), are taught to children, might make the agent a victim themselves. We might imagine that other false moral views were prevalent in this environment and that the agent might have suffered from resulting wrongdoing against themselves, inviting compassion and empathy rather than inclinations to blame or punish. Second, describing agents from specific cultures or societies might lead to a reduction of blame ascriptions out of relativist considerations. This point is especially relevant for empirical studies given the prevalence of relativist intuitions in lay people (Pölzler & Wright, 2020). From this perspective, learning that the agent acted in line with the dominant moral views of their immediate environment might reduce perceived act wrongness which in turn might reduce ascriptions of blame. The agent would not be excused for their wrongdoing but rather their action would not be considered a wrongdoing at all (or at least not to the same degree as otherwise).

In the context of blame and moral ignorance, a related argument has recently been made by Weatherson (2019; see also Fricker, 2010). Weatherson argues that blame might require a common moral ground between the blaming person and the wrongdoer, and that this foundation might be insufficient in cases where the agent comes from a morally compromising environment. The reduced ascriptions of blameworthiness for such an agent would thereby be the result of them being exempt from blame and not of them being excused by their ignorance.⁵

Further studies are required to differentiate between effects of moral ignorance resulting from morally compromising environments and effects of the environment itself. Our findings either suggest that highly biased moral evidence is necessary for agents to be excused by their moral ignorance, or that the blame reduction identified by Faraci and Shoemaker was mostly due to environmental differences or the agents thinking their wrongdoing was morally required, and not due to the moral ignorance itself.

Moral Difficulty

For moral difficulty, the mediation analyses in our study revealed a puzzling effect. First, agents were perceived as substantially less blameworthy for moral ignorance regarding hard cases in comparison to easy cases. This led to an indirect blame-reducing effect of moral difficulty regarding the wrongdoing itself, which was in line with our hypotheses. However, this indirect effect was suppressed by a direct blame-enhancing effect of high moral difficulty on act blameworthiness. After controlling for belief blameworthiness, high moral difficulty was associated with a pre-post-increase of act blameworthiness. This latter effect contradicted our theoretical assumptions. A possible explanation might be that while moral difficulty reduces ascriptions of belief blameworthiness, it also erodes the link between belief blameworthiness and changes in act blameworthiness. This would assume that in cases of high moral difficulty, belief blameworthiness would not be associated with changes in act blameworthiness.

⁵ In contrast, Mason (2019) argues that blame might be used to establish a common moral ground: "Sometimes, of course, we bring people into the philosophical community by holding them to its standards even when we know they will not quite get it. The same is true of morality. We blame proleptically in order to bring people in." (p. 157).

Our operationalization of moral difficulty seems worth discussing. We measured moral difficulty as the individual participants' uncertainty regarding their moral evaluation of the respective action. It is widely assumed that difficult moral cases can lead to uncertainty regarding their evaluation (Kortenkamp & Moore, 2014; MacAskill et al., 2020). We therefore assumed that the subjective experience of moral uncertainty would be associated with the participant's perception of how difficult the respective moral issue is. Our measurement of moral difficulty was thus focused on moral difficulty as a characteristic of the respective moral issue, as it was perceived by the individual participant. This is in line with some arguments regarding the role of difficulty for moral ignorance (e.g., FitzPatrick, 2008; Harman, 2022). However, in other philosophical arguments concerning moral ignorance, moral difficulty has been understood as how difficult it would have been for the *agent*, given their epistemic state, to arrive at the right moral view (e.g., Faraci & Shoemaker, 2010, 2014; Hartford, 2022). Even though this latter notion of difficulty might be strongly associated with how blameworthy the agent is for their ignorance, which we had also captured in our study, both constructs should not be equated and it might prove fruitful to separately capture this perspective on moral difficulty in future studies. Additionally, it might be that participants are highly uncertain about some moral issues of which they assume that they are easy for other people. This would not have been captured in our study.

Findings Regarding the Control Condition

We were surprised to find a small but significant blame reduction in the control trials, where it was described that the moral beliefs of the agent are not known. This blame reduction might have been a mere retest effect that occurs when participants are asked to judge the same behavior twice. The emotional response to the described wrongdoing might be more pronounced the first time a participant reads it, so that by the second time, moral emotions and resulting blame ascriptions might be more moderate (Ginther et al., 2022; Molho et al., 2020; Nelissen & Zeelenberg, 2009). The blame-reduction in the control trials might also have been the result of perspective-taking. Just mentioning that the agent has moral beliefs, even if they are not known to the participants, might have invited a consideration of contrasting moral views, leading to a more moderate blame judgment (Mata, 2019). However, we would have expected that such an effect also manifests in changes in act wrongness, which was not the case in our study. As far as we are aware, our study was the first to use a control condition in a blameupdating paradigm, and our findings suggest the necessity of including one. In further research, the suitability of different control conditions for such paradigms needs to be tested.

Limitations

One might object to our choice of the blame-updating task. The blame-updating task is associated with a psychological understanding of blame as being highly socially regulated, and thus accurate, in contrast to being motivated (Monroe & Malle, 2019). It stands in contrast to other psychological accounts of blame that assume that people are biased towards holding people accountable for negative actions (e.g., Alicke, 2000; Clark et al., 2021; Ditto et al., 2009; Mazzocco et al., 2004). For example, Clark et al. (2021) found evidence of motivated free will attributions which serve to hold wrongdoers responsible, as they compared ascriptions of free will to agents who committed a wrongdoing versus a morally neutral action. Similarly, it might be that participants in our study were motivated to ascribe belief blameworthiness to be able to hold the agent responsible for their action.

If the motivated-blame account is right, our findings show that people tend to hold morally ignorant wrongdoers responsible, but exculpatory effects might be identified in crosssectional comparisons.⁶ If the socially-regulated account of blame is right, as evidence presented by Monroe and Malle (2019) suggests, our findings show that people hold morally ignorant wrongdoers responsible and that they think that this is socially the right thing to do.

⁶ Note that in Monroe and Malle (2019), the blame-updating task led to similar results as a cross-sectional assessment.

The final issue we would like to focus on is the mode of stimulus categorization and the resulting low ratings of moral difficulty. We needed to ensure that the wrongdoings we would present in our study would actually be perceived as wrongdoings by the study participants (Skitka et al., 2021). We therefore asked participants themselves to select actions that they see as morally wrong. The available actions to choose from were presented all at once, and participants were aware that their aim is to choose five morally wrong actions. This mode of presentation might have led to a selection bias where participants mostly chose actions whose moral wrongness was highly salient. We assume that these were issues of rather low moral difficulty, which was reflected in the low mean of moral difficulty ratings. Incorporating issues with a large range of moral difficulty might lead to even stronger associations between moral difficulty and belief blameworthiness and it might help clarify the intricate associations between moral difficulty and change in act blameworthiness. An alternative mode of stimulus categorization might be to have the moral difficulty of actions rated one by one, so that the researchers can then select the most informative stimuli for the blame-updating task.

Further Implications

The debate regarding exculpatory moral ignorance partially parallels that regarding decision-making under moral uncertainty (Harman, 2022; MacAskill et al., 2020). Both debates take place in the larger context of the conflict between normative internalism and normative externalism (Weatherson, 2019). Internalists hold that there are normative standards relative to the agent's moral epistemic state. Put simply, internalism implies that an agent should act in line with the moral views they hold or with the moral evidence that is available to them. Two implications of internalism are thus that certain kinds of moral ignorance are exculpatory and that an agent's credences in different moral views are relevant to what they should do. Externalism rejects all of these claims. The key claim of externalism is that the only relevant standard of evaluation is what morality *actually* demands, regardless of the agent's moral epistemic state (Harman, 2022; Weatherson, 2019). Consequently, externalism rejects the idea

of exculpatory moral ignorance and it rejects the search for meta-normative theories of decision-making under moral uncertainty (uncertaintism) (for further discussions of exculpatory moral ignorance in the context of moral uncertainty, see Geyer, 2018; Guerrero, 2007; Harman, 2015; Rosenthal, 2021; Sepielli, 2017).

Our findings can be interpreted as giving support for the externalist view. If moral ignorance is interpreted as a limiting case of moral uncertainty (Harman, 2015), the lack of an exculpatory effect of moral ignorance suggests that lay people do not link an agent's blameworthiness to the agent's moral beliefs, so that both moral ignorance and moral uncertainty are irrelevant to the evaluation of the agent. Together with recent findings that lay people hold metaethical views that are not in line with the assumptions of uncertaintism (Theisen, 2023), this would cast further doubt regarding the relevance and the intuitive appeal of the uncertaintism project to a lay public. However, our findings also suggested that wrongdoers might be excused if they were completely blameless for the moral epistemic state that their wrongdoing arose from. In our view, it might be easier for an agent to be blamelessly uncertainty (see also Biebel, 2023). From this perspective, the lack of exculpatory moral ignorance might not directly translate to the uncertaintism debate.

Conclusion

Philosophical discussions of moral responsibility have the potential to fundamentally change who we see as liable for blame or praise. The idea of exculpatory moral ignorance suggests that wrongdoers, who we would otherwise have held responsible, might not be as blameworthy as initially assumed. The here presented findings suggest that lay people by and large do not accept the exculpatory potential of moral ignorance. However, we also found some suggestive support for the intuition that a morally ignorant agent is blameless for a wrongdoing if they are fully blameless for their ignorance. Our findings stand in contrast to previous studies suggesting that morally compromising environments can lead to strong reductions in perceived blameworthiness. Future research needs to address which processes are at play here and how

much of such effects can actually be attributed to moral ignorance.

References

- Alicke, M. D. (2000). Culpable control and the psychology of blame. *Psychological Bulletin*, *126*(4), 556–574.
- Arpaly, N. (2003). Unprincipled virtue: An inquiry into moral agency. Oxford University Press.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using Ime4. *Journal of Statistical Software*, *67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01
- Biebel, N. (2023). Reasonable standards and exculpating moral ignorance. *Philosophical Studies*. Advance online publication. https://doi.org/10.1007/s11098-023-02054-x
- Bradford, G. (2017). Hard to know. In P. Robichaud & J. W. Wieland (Eds.), *Responsibility: The epistemic condition* (pp. 180–198). Oxford University Press.
- Clark, C. J., Winegard, B. M., & Shariff, A. F. (2021). Motivated free will belief: The theory, new (preregistered) studies, and three meta-analyses. *Journal of Experimental Psychology: General*, 150(7), e22-e47. https://doi.org/10.1037/xge0000993
- Ditto, P. H., Pizarro, D. A., & Tannenbaum, D. (2009). Motivated moral reasoning. In D. M. Bartels, C. W. Bauman, L. J. Skitka, & D. L. Medin (Eds.), *The psychology of learning and motivation* (pp. 307–338). Academic Press.
- Faraci, D., & Shoemaker, D. (2010). Insanity, deep selves, and moral responsibility: The case of JoJo. *Review of Philosophy and Psychology*, 1(3), 319–332.
- Faraci, D., & Shoemaker, D. (2014). Huck vs. JoJo: Moral ignorance and the (a)symmetry of praise and blame. In J. Knobe, T. Lombrozo, & S. Nichols (Eds.), Oxford studies in experimental philosophy. Oxford studies in experimental philosophy (pp. 7–27). Oxford University Press.
- Fiedler, K., Schott, M., & Meiser, T. (2011). What mediation analysis can (not) do. *Journal of Experimental Social Psychology*, 47(6), 1231–1236.

https://doi.org/10.1016/j.jesp.2011.05.007

FitzPatrick, W. J. (2008). Moral responsibility and normative ignorance: Answering a new skeptical challenge. *Ethics*, *118*(4), 589–613. https://doi.org/10.1086/589532

Fricker, M. (2010). The relativism of blame and Williams's relativism of distance. *Aristotelian Society Supplementary Volume*, *84*(1), 151–177. https://doi.org/10.1111/j.1467-8349.2010.00190.x

Geyer, J. (2018). Moral uncertainty and moral culpability. *Utilitas, 30*(4), 399–416. https://doi.org/10.1017/S0953820818000080

Ginther, M. R., Hartsough, L. E. S., & Marois, R. (2022). Moral outrage drives the interaction of harm and culpable intent in third-party punishment decisions. *Emotion*, 22(4), 795–804. https://doi.org/10.1037/emo0000950

Guerrero, A. A. (2007). Don't know, don't kill: Moral ignorance, culpability, and caution. *Philosophical Studies*, *136*(1), 59–97. https://doi.org/10.1007/s11098-007-9143-7

Guerrero, A. A. (2017). Intellectual difficulty and moral responsibility. In P. Robichaud & J. W. Wieland (Eds.), *Responsibility: The epistemic condition* (pp. 199–218). Oxford University Press.

Harman, E. (2011). Does moral ignorance exculpate? *Ratio*, *24*(4), 443–468. https://doi.org/10.1111/j.1467-9329.2011.00511.x

- Harman, E. (2015). The irrelevance of moral uncertainty. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 10, pp. 53–79). Oxford University Press.
- Harman, E. (2022). Ethics is hard! What follows? In D. K. Nelkin & D. Pereboom (Eds.), *The Oxford handbook of moral responsibility* (pp. 327–347). Oxford University Press.

Hartford, A. (2019). Moral and factual ignorance: A quality of will parity. *Ethical Theory and Moral Practice*, *22*(5), 1087–1102. https://doi.org/10.1007/s10677-019-10043-5

Hartford, A. (2022). Difficulty & quality of will: implications for moral ignorance. *Philosophical Explorations*, *25*(2), 141–158. https://doi.org/10.1080/13869795.2022.2042585

Jennings, M. A., & Cribbie, R. A. (2016). Comparing pre-post change across groups: Guidelines for choosing between difference scores, ANCOVA, and residual change scores. *Journal of Data Science*, *14*, 205–230.

Kievit, R. A., Brandmaier, A. M., Ziegler, G., van Harmelen, A.-L., Mooij, S. M. M. de,
Moutoussis, M., Goodyer, I. M., Bullmore, E., Jones, P. B., Fonagy, P., Lindenberger, U., &
Dolan, R. J. (2018). Developmental cognitive neuroscience using latent change score
models: A tutorial and applications. *Developmental Cognitive Neuroscience*, *33*, 99–117.
https://doi.org/10.1016/j.dcn.2017.11.007

- Kirfel, L., & Hannikainen, I. R. (2023). Why blame the ostrich? Understanding culpability for willful ignorance. In S. Magen & K. Prochownik (Eds.), Advances in experimental philosophy of law (pp. 75–98). Bloomsbury Academic.
- Kissinger-Knox, A., Aragon, P., & Mizrahi, M. (2018). Does non-moral ignorance exculpate?
 Situational awareness and attributions of blame and forgiveness. *Acta Analytica*, 33, 161–179.
- Kortenkamp, K. V., & Moore, C. F. (2014). Ethics under uncertainty: The morality and appropriateness of utilitarianism when outcomes are uncertain. *American Journal of Psychology*, *127*(3), 367. https://doi.org/10.5406/amerjpsyc.127.3.0367
- MacAskill, W., Bykvist, K., & Ord, T. (Eds.). (2020). Moral uncertainty. Oxford University Press.
- Makins, N. (2021). Attitudinal ambivalence: Moral uncertainty for non-cognitivists. *Australasian Journal of Philosophy*, 1–15. https://doi.org/10.1080/00048402.2021.1908380
- Mason, E. (2019). *Ways to be blameworthy: Rightness, wrongness, and responsibility*. Oxford University Press.
- Mata, A. (2019). Social metacognition in moral judgment: Decisional conflict promotes perspective taking. *Journal of Personality and Social Psychology*, *117*(6), 1061–1082. https://doi.org/10.1037/pspa0000170
- Mazzocco, P. J., Alicke, M. D., & Davis, T. L. (2004). On the robustness of outcome bias: No constraint by prior culpability. *Basic and Applied Social Psychology*, *26*(2-3), 131–146. https://doi.org/10.1080/01973533.2004.9646401
- McNeish, D., Stapleton, L. M., & Silverman, R. D. (2017). On the unnecessary ubiquity of hierarchical linear modeling. *Psychological Methods*, 22(1), 114–140. https://doi.org/10.1037/met0000078
- Molho, C., Tybur, J. M., van Lange, P. A. M., & Balliet, D. (2020). Direct and indirect punishment of norm violations in daily life. *Nature Communications*, *11*(1), 3432. https://doi.org/10.1038/s41467-020-17286-2
- Monroe, A. E., & Malle, B. F. (2019). People systematically update moral judgments of blame. Journal of Personality and Social Psychology, 116(2), 215–236. https://doi.org/10.1037/pspa0000137
- Nelissen, R. M. A., & Zeelenberg, M. (2009). Moral emotions as determinants of third-party punishment: Anger, guilt, and the functions of altruistic sanctions. *Judgment and Decision Making*, *4*(7), 543–553. https://doi.org/10.1017/s1930297500001121
- Pölzler, T., & Wright, J. C. (2020). Anti-realist pluralism: A new approach to folk metaethics.
 Review of Philosophy and Psychology, *11*(1), 53–82. https://doi.org/10.1007/s13164-019-00447-8
- R Core Team. (2022). *R: A language and environment for statistical computing* (Version 4.2.2) [Computer software]. Vienna, Austria. https://www.R-project.org/
- Robichaud, P., & Wieland, J. W. (Eds.). (2017). *Responsibility: The epistemic condition*. Oxford University Press.
- Rosen, G. (2003). Culpability and ignorance.
- Rosen, G. (2004). Skepticism about moral responsibility. *Philosophical Perspectives*, *18*(Ethics), 295–313.

- Rosenthal, C. (2021). What decision theory can't tell us about moral uncertainty. *Philosophical Studies*, *178*(10), 3085–3105. https://doi.org/10.1007/s11098-020-01571-3
- Rosseel, Y. (2012). lavaan : An R Package for Structural Equation Modeling. *Journal of Statistical Software, 48*(2). https://doi.org/10.18637/jss.v048.i02
- Schwartz, S. A., & Inbar, Y. (2023). Is it good to feel bad about littering? Conflict between moral beliefs and behaviors for everyday transgressions. *Cognition*, *236*, 105437. https://doi.org/10.1016/j.cognition.2023.105437
- Sepielli, A. (2017). How moral uncertaintism can be both true and interesting. In M. C. Timmons
 (Ed.), Oxford studies in normative ethics (pp. 98–116). Oxford University Press.
 https://doi.org/10.1093/oso/9780198808930.003.0006
- Skitka, L. J., Hanson, B. E., Morgan, G. S., & Wisneski, D. C. (2021). The psychology of moral conviction. *Annual Review of Psychology*, 72, 347–366. https://doi.org/10.1146/annurev-psych-063020-030612
- Sliwa, P. (2020). Excuse without exculpation. In P. Sliwa (Ed.), *Oxford studies in metaethics Volume 15* (pp. 72–95). Oxford University Press.

https://doi.org/10.1093/oso/9780198859512.003.0004

Smith, H. M. (1983). Culpable ignorance. *The Philosophical Review*, *92*(4), 543–571.

- Steyer, R., Eid, M., & Schwenkmezger, P. (1997). Modeling true intraindividual change: True change as a latent variable. *Methods of Psychological Research Online*, *2*(1), 21–33.
- Theisen, M. (2023). Metaethical intuitions in lay concepts of normative uncertainty. *Philosophical Psychology*. Advance online publication.

https://doi.org/10.1080/09515089.2023.2174422

- Weatherson, B. (2019). Normative externalism. Oxford University Press.
- Wieland, J. W. (2017). What's special about moral ignorance? Ratio, 30(2), 149–164.

https://doi.org/10.1111/rati.12127

Wolf, S. (1987). Sanity and the metaphysics of responsibility. In F. Schoeman (Ed.), *Responsibility, character, and the emotions: New essays in moral psychology* (pp. 363–373). Cambridge University Press.

Zimmerman, M. J. (1997). Moral responsibility and ignorance. *Ethics*, *107*(3), 41–426.

Act description	Personal act description	Times selected as immoral by participants (percentage of participants)
Eating meat	Nick often eats meat.	12 (4.8%)
Illegally downloading movies	Kaja often downloads movies illegally.	8 (3.2%)
Using "cripple" as an insult	Milan often uses "cripple" as an insult.	69 (27.5%)
Avoiding giving money to the homeless	Philipp often avoids giving money to the homeless.	1 (0.4%)
Slacking off while at work	Martin is often slacking off while at work.	6 (2.4%)
Buying things that you do not need	Thorsten often buys things that he does not need.	4 (1.6%)
"Checking out" strangers on the street	Jan often checks out strangers on the street.	1 (0.4%)
Visiting a senior residence without a mask	Lea often visits a senior residence without a mask.	18 (7.2%)
Voting for fascist parties	Rebecca often votes for fascist parties.	147 (58.6%)
Ordering clothing that you expect to return	Cem often orders clothing that he expects to return.	6 (2.4%)
Trying to convert people to your own faith	Fabienne often tries to convert people to her faith.	48 (19.1%)
Having one-night stands with strangers	Kim often has one-night stands with strangers.	6 (2.4%)
Consuming cannabis	Hannes often consumes cannabis.	1 (0.4%)
Physically punishing your children	Tim often punishes his children physically.	176 (70.1%)
Protesting against abortions in front of hospitals	Marie often protests against abortions in front of hospitals.	69 (27.5%)

abortions.

malls.

assisted suicide.

As a doctor, Jasmin often performs

As a doctor, Tina often performs

Steven often steals in shopping

7 (2.8%)

7 (2.8%)

71 (28.3%)

Performing abortions as a doctor

Performing assisted suicide as a

Stealing in shopping malls

doctor

List of Actions That Could Be Selected as Stimulus Material by the Participants

Act description	Personal act description	Times selected as immoral by participants (percentage of participants)
Flying to Milan over the weekend	Lars often flies to Milan over the weekend.	8 (3.2%)
Blocking traffic as a form of climate activism	Lara often blocks traffic as a form of climate activism.	49 (19.5%)
Demonstrating against animal testing	Sandra often demonstrates against animal testing.	-
Carrying out painful animal experiments as a researcher	As a researcher, Emil often carries out painful animal experiments.	103 (41%)
Sharing photos of your child unpixellated on social media	Barbara often shares photos of her child unpixellated on social media.	50 (19.9%)
Cheating in exams	Kerstin often cheats in exams.	16 (6.4%)
Posting hate comments about politicians	Ralf often posts hate comments about politicians.	52 (20.7%)
Reporting the neighbors to the authorities for minor offenses	Jenny often reports her neighbors to the authorities for minor offenses.	64 (25.5%)
Refusing to tip in the restaurant	Maya often refuses to tip in the restaurant.	3 (1.2%)
Investing in shares of the defense industry	Otto often invests in shares of the defense industry.	29 (11.6%)
Calling in sick without actually being sick	Theresa often calls in sick without actually being sick.	29 (11.6%)
Failing to do one's part in a group work	Ben often fails to do his part in a group work.	32 (12.7%)
Selling nude pictures of oneself online	Lisa often sells nude pictures of herself online.	16 (6.4%)
Watching porn	Finn often watches porn.	5 (2%)
Going to work despite cold symptoms	Özge often goes to work despite cold symptoms.	31 (12.4%)
Talking loudly to one's friends during movies at the cinema	Paul often talks loudly to his friends during movies at the cinema.	51 (20.3%)
Smuggling sweets into the cinema	Antonia often smuggles sweets into the cinema.	-
Trying to convince others to go vegan	Serkan often tries to convince others to go vegan.	19 (7.6%)

Act description	Personal act description	Times selected as immoral by participants (percentage of participants)
Cancelling meetings at short notice if one does not feel like it	Jessica often cancels meetings at short notice if she does not feel like it.	15 (6%)
Giving unsolicited advice	Sebastian often gives unsolicited advice.	1 (0.4%)
Using public transport without buying a ticket	Melissa often uses public transport without buying a ticket.	25 (10%)

Exemplary Manipulations of Procedural Obligations and Moral Evidence in the Moral Ignorance

Trials

Procedural	Moral	Ctinulus tout
obligations	evidence	Stimulus text
Fulfilled	Biased	Nick often eats meat. Nick has thought very carefully about the moral aspects of this action. He has also critically questioned his own opinions and sentiments on the subject. However, he was only confronted with arguments and experiences that indicated that it was morally acceptable to eat meat. He did not realize that his considerations were one-sided as a result and ultimately concluded that it was morally acceptable to eat meat.
	Unbiased	Nick often eats meat. Nick has thought very carefully about the moral aspects of this action. He has also critically questioned his own opinions and sentiments on the subject. In doing so, he was also confronted with arguments and experiences that indicated that it might be morally wrong to eat meat. After very careful consideration, however, the other arguments prevailed for him and so he ultimately concluded that it is morally acceptable to eat meat.
Violated	Biased	Nick often eats meat. Nick has thought only superficially about the moral aspects of this action. He largely followed his first intuition on the subject without critically questioning his own opinions and sentiments. In addition, he was only confronted with arguments and experiences that indicated that it was morally acceptable to eat meat. Thus, he concluded that it was morally acceptable to eat meat.
	Unbiased	Nick often eats meat. Nick has thought only superficially about the moral aspects of this action. He largely followed his first intuition on the subject without critically questioning his own opinions and sentiments. He was admittedly also confronted with arguments and experiences that indicated that it might be morally wrong to eat meat. However, he did not consider these further in his moral evaluation of the issue. Thus, he concluded that it is morally acceptable to eat meat.

Model	Syntax	σ_{subj}	σ_{act}	σ_{res}	R ² _m	R ² c	AIC
1a	wrongness_diff ~ ignorance + (1 person) + (1 action)	2.49	0.00ª	7.84	0.0%	9.8%	8839.5
1b	wrongness_diff ~ ignorance + (1 person)	2.49	-	7.84	0.0%	9.2%	8837.5
1c	wrongness_diff ~ 0 + ignorance + (1 person)	2.49	-	7.84	0.0%	9.2%	8837.5
2a	actblame_diff ~ ignorance + (1 person) + (1 action)	3.27	0.98	9.62	0.3%	11.5%	9375.1
2b	actblame_diff ~ 0 + ignorance + (1 person) + (1 action)	3.27	0.98	9.62	0.3%	11.5%	9375.1
2c	actblame_diff ~ ignorance + difficulty + (1 person) + (1 action)	3.29	0.94	9.62	0.3%	11.5%	9382.6
2d	actblame_diff ~ ignorance * difficulty + (1 person) + (1 action)	3.29	0.93	9.61	0.6%	11.8%	9386.0
3	actblame_diff ~ condition + (1 person) + (1 action)	3.38	1.05	9.44	0.3%	15.1%	9336.4
4a	actblame_diff ~ procedural * evidence + difficulty + (1 person) + (1 action)	3.64	1.57	9.78	3.4%	17.1%	7566.0
4b	actblame_diff ~ procedural * evidence + (1 person) + (1 action)	3.63	1.61	9.78	3.3%	17.0%	7559.3
5	actblame_diff ~ beliefblame + (1 person) + (1 action)	3.15	2.02	9.46	11.9%	23.8%	7489.9
6a	beliefblame ~ procedural * evidence + difficulty + (1 person) + (1 action)	9.76	4.32	11.25	8.6%	51.9%	8108.4
6b	beliefblame ~ procedural + evidence + difficulty + (1 person) + (1 action)	9.77	4.26	11.25	8.5%	51.8%	8110.1
6c	beliefblame ~ evidence + difficulty + (1 person) + (1 action)	9.77	4.26	11.25	8.5%	51.8%	8109.5

Overview of Linear Mixed Models Fitted for This Study

Note. Models 1a-3b were fitted on the full dataset including control and moral ignorance trials

(1255 observations). Models 4a-6b were fitted on a subset of the data including only moral

ignorance trials (1004 observations). $\sigma_{subj:}$ standard deviation between participants; σ_{act} : standard deviation between moral issues; σ_{res} : residual standard deviation; R^2_m : proportion of variance explained by fixed effects; R^2_c : proportion of variance explained by fixed and random effects; *AIC*: Akaike information criterion.

^a Singular fit.

Descriptive Statistics and Correlations for Each of the Assessed Variables

						Co	orrelation	S		
	Variable	М	SD	1	2	3	4	5	6	7
1	Act blame- worthiness (Pre)	89.3	13.6	-						
2	Act blame- worthiness (Post)	89.0	14.6	.739***	-					
3	Change in act blame- worthiness	-0.3	10.2	281***	.439***	-				
4	Belief blame- worthiness	87.0	16.5	.651***	.848***	.317*	-			
5	Act wrongness (Pre)	88.5	15.7	.642***	.547***	053	.570***	-		
6	Act wrongness (Post)	88.6	16.1	.630***	.611***	.030	.635***	.867***	-	
7	Change in act wrongness	0.1	8.2	.047	.157***	.161***	.163***	203****	.312***	-
8	Moral difficulty	11.0	17.0	435***	390***	.024	373***	585	567	.003

Note. The theoretical range of each variable was 0 - 100, except for the change scores which had a theoretical range of -100 - 100. Correlations involving belief blameworthiness refer to moral ignorance trials only.

*****p* < .001.

Fixed Effect Estimates for the Model Predicting Changes in Act Blameworthiness From Procedural

Variable	b	95% Cl.lower	95% Cl.upper	t (df)	р	$R^2_{partial}$
Intercept	1.15	-0.34	2.62	1.522 (118.16)	.131	-
Procedural obligations (fulfilled)	0.91	0.81	2.64	1.037 (747.157)	.300	0.2%
Moral evidence (biased)	-1.94	-3.65	-0.21	-2.21 (741.099)	.027	0.7%
Fulfilled x biased	-3.13	-5.56	-0.70	-2.522 (742.222)	.012	0.9%

Obligations, Moral Evidence, and Their Interaction (Model 4b)

Fixed Effect Estimates for the Model Predicting Belief Blameworthiness From Procedural

Variable	b	95% Cl.lower	95% Cl.upper	t (df)	p	$R^2_{\rm partial}$
Intercept	89.82	87.14	92.42	67.754 (72.841)	< .001	-
Procedural obligations (fulfilled)	1.10	-0.92	3.10	1.070 (728.219)	.285	0.1%
Moral evidence (biased)	-2.76	-4.75	-0.77	-2.717 (717.510)	.007	0.8%
Fulfilled x biased	-1.54	-4.35	1.29	-1.067 (720.182)	.286	0.1%
Moral difficulty	-0.26	-0.32	-0.21	-9.672 (961.703)	< .001	13.7%

Obligations, Moral Evidence, Their Interaction, and Moral Difficulty (Model 6a)

Standardized Estimates From Manifest Path Analysis and Latent Change Score Modeling for Indirect and Total Effects of Procedural Obligations, Moral

Evidence, Their Interaction, and Moral Difficulty Through Belief Blameworthiness on Changes of Act Blameworthiness

			+covince	offocto			Totol	ffooto	
			Indirect	errects			I OTAI (errects	
		Point	95%	95%		Point	95%	95%	
Model	Variable	estimate	CI.lower	Cl.upper	d	estimate	CI.lower	Cl.upper	d
Path analysis									
	Procedural obligations (fulfilled)	0.010	-0.013	0.034	.403	0.045	-0.028	0.118	.222
	Moral evidence (biased)	-0.033	-0.058	-0.008	.010	-0.086	-0.163	-0.008	.030
	Fulfilled x biased	-0.013	-0.041	0.015	.369	-0.127	-0.220	-0.034	.007
	Moral difficulty	-0.140	-0.198	-0.082	< .001	0.039	-0.047	0.125	.374
LCSM									
	Procedural obligations (fulfilled)	0.017	-0.016	0.051	.310	0.054	-0.036	0.144	.240
	Moral evidence (biased)	-0.042	-0.075	-0.008	.015	-0.096	-0.193	0.000	.050
	Fulfilled x biased	-0.021	-0.060	0.018	0.291	-0.146	-0.262	-0.029	.014
	Moral difficulty	-0.177	-0.254	-0.101	< .001	0.057	-0.062	0.175	.350

Note. LCSM: latent change score modeling.



Theoretical Structural Equation Model That Was Fitted to the Data

Note. Theoretical structural equation model that was fitted to the data. It contains a latent change variable, as well as direct and mediated effects of fulfilled procedural obligations, biased moral evidence, their interaction, and moral difficulty on latent change. It further contains covariances between pre-measurements of act blameworthiness and the predictor variables (including the mediator), as well as covariances between similar items across different time points (pre vs. post) and different objects of blame (act vs. belief). Intercepts are omitted from the diagram for clarity. Numbers indicate fixed parameters. p2 and p3 describe parameters fixed to equality. Resp.: responsibility. Pre: pre-measurement of act blameworthiness. Post: post-measurement of act blameworthiness. Belief: belief blameworthiness. Change: true change in act blameworthiness between pre- and post-measurement.

Pre-Post Changes in Act Blameworthiness by Experimental Condition (Model 3)



Note. Possible values ranged from -100 to 100. Positive values indicate increased blame, negative values indicate a blame reduction. Bars indicate estimated marginal means for each condition including a 95% confidence interval. Dunnett-corrected *p*-values and Cohen's *d* for each comparison against the control condition are given. As a reference, the outmost right bar indicates the blame reduction Monroe & Malle (2019) found in their first study for acting unintentionally.

Marginal Effects for Predicting Changes in Act Blameworthiness From Belief Blameworthiness



Including 95% Confidence Band (Model 5)



Marginal Effects for Predicting Belief Blameworthiness (Model 6a)

Note. A) Belief blameworthiness by experimental condition. Bars indicate estimated marginal means for each condition including a 95% confidence interval. Marginal means were estimated for an average moral difficulty of 11. B) Predicted values of belief blameworthiness from moral difficulty including 95% confidence band.

Manifest Path Analysis



Note. Mediation effects of fulfilled procedural obligations, biased moral evidence, their interaction, and moral difficulty via belief blameworthiness on changes in act blameworthiness. Standardized path estimates (*p*-values) are provided. Solid (dotted) paths are (not) significant at a 5% significance level.



Extract of Path Estimates From the Fitted Full Structural Equation Model (see Figure 1)

Note. Mediation effects of fulfilled procedural obligations, biased moral evidence, their interaction, and moral difficulty via latent belief blameworthiness on latent change in act blameworthiness. Standardized path estimates (*p*-values) are provided. Solid (dotted) paths are (not) significant at a 5% significance level.

Appendix: Supplementary Material

Code S1

```
## Measurement models
# Define latent constructs
PRE =~ 1*pre1 + p2*pre2 + p3*pre3
                                        # PRE: Act blameworthiness pre
POST =~ 1*post1 + p2*post2 + p3*post3
                                        # POST: Act blameworthiness post
                                        # MED: Belief blameworthiness
MED =~ 1*med1 +
                  med2 +
                             med3
# Intercepts of latent variables
PRE ~ 0*1
                                         # Fix intercept of PRE to 0
POST ~ 0*1
                                         # Fix intercept of POST to 0
MED ~ 1
                                         # Estimate intercept for MED
# Variance of latent variables
PRE ~~ PRE
                                         # Estimate variance in PRE
POST ~~ 0*POST
                                         # Fix variance in POST to 0
MED ~~ MED
                                         # Estimate variance in MED
# Covariance of latent variables
PRE ~~ MED
                                         # Estimate covariance between PRE and MED
# Intercepts of indicators
pre1 ~ 0*1
pre2 ~ 0*1
pre3 ~ 0*1
post1 ~ 0*1
post2 ~ 0*1
post3 ~ 0*1
med1 ~ 0*1
med2 ~ 0*1
med3 ~ 0*1
# Residual variances of indicators
pre1 ~~ pre1
pre2 ~~ pre2
pre3 ~~ pre3
post1 ~~ post1
post2 ~~ post2
post3 ~~ post3
med1 ~~ med1
med2 ~~ med2
med3 ~~ med3
# Covariances between corresponding indicators (1: responsibility, 2: blame, 3:
reproach)
pre1 ~~ post1
pre2 ~~ post2
pre3 ~~ post3
pre1 ~~ med1
pre2 ~~ med2
pre3 ~~ med3
post1 ~~ med1
post2 ~~ med2
post3 ~~ med3
# Define latent change variable
CHANGE =~ 1*POST
                                         # CHANGE: POST - PRE
POST ~ 1*PRE
                                         # Fix path from PRE to POST to 1
CHANGE ~~ PRE
                                         # Estimate covariance between PRE and CHANGE
CHANGE ~ 1
                                         # Estimate intercept for CHANGE
CHANGE ~~ CHANGE
                                         # Estimate variance in CHANGE
```

```
## Predictors & mediation
# Fulfilled procedural obligations
procedural ~ 0.5*1
procedural ~~ 0.25*procedural
# Biased moral evidence
epistemic ~ 0.5*1
epistemic ~~ 0.25*epistemic
# Interaction of fulfilled procedural obligations and biased moral evidence
fulfilled biased ~ -0.25*1
fulfilled biased ~~ 0.0625*fulfilled biased
fulfilled_biased ~ 0.5*procedural + 0.5*epistemic
# Moral difficulty
difficulty ~ 0*1
difficulty ~~ 1*difficulty
# Regression of CHANGE on predictors and MED
CHANGE ~ b1*MED + c1*procedural + c2*epistemic + c3*fulfilled_biased + c4*difficulty
# Regression of MED on predictors
MED ~ a1*procedural + a2*epistemic + a3*fulfilled_biased + a4*difficulty
# Covariance between PRE and predictors
PRE ~~ difficulty + procedural + epistemic + fulfilled_biased
# Mediation terms
indirect_procedural := a1 * b1
indirect_epistemic := a2 * b1
indirect_interaction := a3 * b1
indirect_difficulty := a4 * b1
total_procedural := c1 + a1 * b1
total_epistemic := c2 + a2 * b1
total interaction := c3 + a3 * b1
total_difficulty := c4 + a4 * b1
```

Table S1

Model Fits for Different Levels of Measurement Invariance Between Pre- and Post-Measurements

of Act Blameworthiness

Model	CElectron	TI Instant	RMSFA	RMSEA _{robust}	SRMR
	Chrobust	, Litobust	TITIO E TODUST	5070 61	3/////
Overall model	1.000	1.005	.000	[.000; .000]	0.000
Pre only model	1.000	1.007	.000	[.000; .000]	0.000
Post only model	1.000	1.004	.000	[.000; .000]	0.000
Configural invariance model	1.000	1.005	.000	[.000; .000]	0.000
Weak invariance model	1.000	0.999	.014	[.000; .072]	0.024
Strong invariance model	1.000	1.000	.000	[.000; .060]	0.024
Strict invariance model	0.975	0.981	.080	[.035; .126]	0.036

Note. Invariance given if difference in CFI_{robust} between parsimonious and less parsimonious

model <= 0.01.

Table S2

Variable	Estimate	SE	Ζ	р
		Factor loadings		
PRE				
pre1	1.00+			
pre2	1.15	0.09	13.44	< .001
pre3	1.05	0.09	12.12	< .001
POST				
post1	1.00+			
post2	1.15	0.09	13.44	< .001
post3	1.05	0.09	12.12	< .001
MED				
med1	1.00+			
med2	1.13	0.08	13.99	< .001
med3	0.98	0.11	8.55	< .001
CHANGE				
POST	1.00+			
		Regression slopes		
POST				
PRE	1.00+			
fulfilled_biased				
procedural	0.50+			
epistemic	0.50+			
CHANGE				
MED	0.29	0.05	5.39	< .001
procedural	0.03	0.04	0.78	.436
epistemic	-0.05	0.05	-1.13	.259
fulfilled_ biased	-0.14	0.07	-2.03	.042
difficulty	0.11	0.03	3.79	< .001
MED				
procedural	0.06	0.06	0.99	.320
epistemic	-0.13	0.05	-2.52	.012

Unstandardized Estimates for the Full Latent Change Score Model

Variable	Estimate	SE	Ζ	p
fulfilled_ biased	-0.08	0.08	-1.03	.302
difficulty	-0.29	0.04	-6.57	< .001
		Intercepts		
pre1	0.00+			
pre2	0.00+			
pre3	0.00+			
post1	0.00+			
post2	0.00+			
post3	0.00+			
med1	0.00+			
med2	0.00+			
med3	0.00+			
difficulty	0.00+			
procedural	0.50+			
epistemic	0.50+			
fulfilled_biased	-0.25+			
		Residual variances		
pre1	0.63	0.12	5.1	< .001
pre2	0.44	0.07	6.38	< .001
pre3	0.42	0.06	7.67	< .001
post1	0.41	0.07	6.09	< .001
post2	0.26	0.05	5.23	< .001
post3	0.36	0.06	6.07	< .001
med1	0.38	0.07	5.4	< .001
med2	0.21	0.05	4.57	< .001
med3	0.37	0.07	5.62	< .001
difficulty	1.00+			
procedural	0.25+			
epistemic	0.25+			
fulfilled_biased	0.06+			
		(Residual) covariances	5	
pre1 w/post1	0.16	0.03	4.71	< .001
pre2 w/post2	0.11	0.04	2.48	.013
pre3 w/post3	0.26	0.05	5.31	< .001

Variable	Estimate	SE	Z	p
pre1 w/med1	0.00	0.03	0.05	.958
pre2 w/med2	-0.03	0.02	-1.22	.222
pre3 w/med3	0.24	0.05	4.91	< .001
post1 w/med1	0.05	0.04	1.17	.241
post2 w/med2	-0.03	0.03	-1.22	.223
post3 w/med3	0.25	0.06	4.31	< .001
PRE w/difficulty	-0.32	0.04	-8.12	< .001
PRE w/procedural	0.02	0.01	1.67	.094
PRE w/epistemic	-0.01	0.01	-0.64	.521
PRE w/fulfilled_ biased	0.00	0.00	-0.23	.819
		Latent intercepts		
PRE	0.00+			
POST	0.00+			
MED	0.06	0.04	1.44	.151
CHANGE	0.04	0.03	1.29	.196
		Latent variances		
PRE	0.44	0.07	6.36	< .001
POST	0.00+			
MED	0.53	0.1	5.24	< .001
CHANGE	0.17	0.03	5.07	< .001
		Latent covariances		
PRE w/MED	0.31	0.05	6.06	< .001
PRE w/CHANGE	-0.13	0.03	-4.76	< .001
		Indirect effects		
procedural	0.02	0.02	1.00	.315
epistemic	-0.04	0.02	-2.29	.022
fulfilled_biased	-0.02	0.02	-1.04	.299
difficulty	-0.08	0.02	-3.96	< .001
		Total effects		
procedural	0.05	0.04	1.16	.245
epistemic	-0.09	0.05	-1.91	.056
fulfilled_biased	-0.16	0.07	-2.35	.019

Variable	Estimate	SE	Ζ	p
difficulty	0.03	0.03	0.94	.347

Note. PRE: act blameworthiness pre. POST: act blameworthiness post. CHANGE: true difference between pre- and post-measurements of act blameworthiness. MED: belief blameworthiness. procedural: fulfilled procedural obligations. epistemic: biased moral evidence. fulfilled_biased: interaction between fulfilled procedural obligations and biased moral evidence. difficulty: moral difficulty.

⁺ Fixed parameter.

Declaration in accordance to § 8 (1) c) and d) of the doctoral degree regulation of the Faculty

FAKULTÄT FÜR VERHALTENS- UND EMPIRISCHE KULTURWISSENSCHAFTEN



Promotionsausschuss der Fakultät für Verhaltens- und Empirische Kulturwissenschaften der Ruprecht-Karls-Universität Heidelberg / Doctoral Committee of the Faculty of Behavioural and Cultural Studies of Heidelberg University

Erklärung gemäß § 8 (1) c) der Promotionsordnung der Universität Heidelberg für die Fakultät für Verhaltens- und Empirische Kulturwissenschaften / Declaration in accordance to § 8 (1) c) of the doctoral degree regulation of Heidelberg University, Faculty of Behavioural and Cultural Studies

Ich erkläre, dass ich die vorgelegte Dissertation selbstständig angefertigt, nur die angegebenen Hilfsmittel benutzt und die Zitate gekennzeichnet habe. / I declare that I have made the submitted dissertation independently, using only the specified tools and have correctly marked all quotations.

Erklärung gemäß § 8 (1) d) der Promotionsordnung der Universität Heidelberg für die Fakultät für Verhaltens- und Empirische Kulturwissenschaften / Declaration in accordance to § 8 (1) d) of the doctoral degree regulation of Heidelberg University, Faculty of Behavioural and Cultural Studies

Ich erkläre, dass ich die vorgelegte Dissertation in dieser oder einer anderen Form nicht anderweitig als Prüfungsarbeit verwendet oder einer anderen Fakultät als Dissertation vorgelegt habe. / I declare that I did not use the submitted dissertation in this or any other form as an examination paper until now and that I did not submit it in another faculty.

Vorname Nachname / First name Family name	Maximilian Johannes Theisen	
Datum / Date	March 15, 2024	
Unterschrift / Signature	Dem Dekanat der Fakultät für Verhaltens- und Empirische Kulturwissenschaften liegt eine unterschriebene Version dieser Erklärung vom 15.03.2024 vor.	