

# **Moral Vagueness**

by

Kaj André Zeller

A Dissertation Submitted to the  
Graduate School of Social Sciences and Humanities  
in Partial Fulfillment of the Requirements for  
the Degree of  
Master of Arts

in

Philosophy



**KOÇ  
ÜNİVERSİTESİ**

28.05.2024

# Moral Vagueness

Koç University

Graduate School of Social Sciences and Humanities

This is to certify that I have examined this copy of a master's thesis by

**Kaj André Zeller**

and have found that it is complete and satisfactory in all respects,  
and that any and all revisions required by the final  
examining committee have been made.

Committee Members:

---

Asst. Prof. David Killoren (Advisor)

---

Prof. Zeynep Direk

---

Assoc. Prof. Bill Wringe

Date: 28.05.2024



*To Bilinç, who clearly is in the extension of the predicate ‘... is loved by Kaj’.*

# **ABSTRACT**

**Moral Vagueness**

**Kaj André Zeller**

**Master of Arts in Philosophy**

**28.05.2024**

Cheryl considers an on-demand abortion after 151 days of gestational age. Is Cheryl morally permitted to do so? Many philosophers believe that cases like this illustrate that our moral vocabulary is vague. However, not everybody's convinced by such alleged examples. In this work, I demonstrate that straightforward arguments for moral vagueness fail due to the characterization problem of vagueness. However, I argue that the question of whether there is moral vagueness must nevertheless be taken seriously. Because it seems that moral predicates are vague, we have reason to believe that moral predicates are vague. This claim is supported by a phenomenological investigation of vagueness-related borderline cases, showing that such cases have a distinct phenomenological profile. In addition to these arguments, I also argue that apparent reasons to dismiss moral vagueness despite the appearance that moral predicates are vague are unconvincing.

## ÖZETÇE

**Ahlaki Muğlaklık**

**Kaj André Zeller**

**Master of Arts in Philosophy**

**28.05.2024**

Cheryl, 151 günlük gebelik süresinden sonra isteğe bağlı kürtaj yapmayı düşünüyor. Cheryl'in bunu yapmasına ahlaki açıdan izin veriliyor mu? Pek çok filozof, bu gibi durumların, ahlaki kelime dağarcığımızın muğlak olduğunu gösterdiğine inanıyor. Ancak, herkes bu tür iddialara ikna olmuş değil. Bu çalışmada, ahlaki muğlaklıkla ilgili basit argümanların, muğlaklığın karakterizasyon sorunu nedeniyle başarısız olduğunu gösteriyorum. Bununla birlikte ahlaki muğlaklığın bulunup bulunmadığı sorusunun yine de ciddiye alınması gerektiğini savunuyorum. Çünkü ahlaki yüklem muğlak görünüyor, ahlaki yüklem muğlak olduğuna inanmak için nedenlerimiz var. Bu iddia, muğlaklıkla ilişkili sınır vakalarına ilişkin fenomenolojik bir araştırmayla destekleniyor ve bu tür vakaların belirgin bir fenomenolojik profile sahip olduğunu gösteriyor. Bu argümanlara ek olarak, ahlaki yüklem muğlak olduğu görünümüne rağmen, ahlaki muğlaklığı göz ardı etmeye yönelik görünürdeki nedenlerin ikna edici olmadığını da savunuyorum.

## **ACKNOWLEDGEMENTS**

Many thanks to David for keeping me going, my fellow grad students at Koç for all your support and gallows humor, and Bilinç, Sazuke, and Arya for everything.



## TABLE OF CONTENTS

List of Figures.....	ix
Chapter 1: Introduction.....	1
Chapter 2: Vagueness in Ethics .....	6
2.1 Vagueness in Communication .....	8
2.2 Moral Indeterminacy.....	10
Chapter 3: Straightforward Arguments for Moral Vagueness.....	12
3.1 The Characterization Problem of Vagueness.....	13
3.1.1 Vagueness as Borderlineness .....	13
3.1.2 Vagueness as Boundarylessness.....	14
3.1.3 Epistemic and Psychological Characterizations .....	16
3.2 Particular Straightforward Arguments.....	17
3.2.1 The Straightforward Argument from Imprecise Gradability.....	18
3.2.2 The Straightforward Argument from Boundarylessness.....	19
3.2.3 The Straightforward Argument from Borderlineness.....	20
3.2.4 The Straightforward Argument from Soriticality/Tolerance.....	21
Chapter 4: Moral Vagueness as a Philosophical Datum.....	22
4.1 Philosophical Data .....	23
4.1.1 The Nature of Data .....	24
4.1.2 What Makes Something a Datum.....	25
4.1.3 Data and Theory .....	26
4.1.4 Why Appeal to Bengson, Cuneo, and Shafer-Landau?.....	28
4.2 Seemings.....	34
4.3 The Seeming That There's Moral Vagueness as a Linguistic Seeming ...	38
4.4 The Phenomenology of Vagueness.....	47
Chapter 5: Handling Moral Vagueness .....	54
5.1 Passing On Moral Vagueness .....	55

5.2	Data-Disablers of the First Kind.....	60
5.2.1	The Kripke-Test.....	62
5.2.2	The Computer-Test.....	63
5.2.3	The Parasitism-Test.....	65
5.2.4	The Moore-Test.....	66
5.3	A Data-Disabler of the Second Kind.....	68
	Chapter 6: Conclusion.....	75
	<b>BIBLIOGRAPHY.....</b>	<b>77</b>



## LIST OF FIGURES

Figure 1: The Semantics of Moral Vagueness.....	42
---	----



## Chapter 1: INTRODUCTION

This thesis began with the question, ‘What should one do when it’s vague what one should do?’ This led to the further question, ‘Are there, actually, situations where it’s vague what one should do?’ which led, in its turn, to research into philosophical methodology as I realized that it is not at all straightforward how to approach the problem of distinguishing the vague from the non-vague.

The central question this thesis investigates is thus how to approach the question of whether there is moral vagueness. To set the stage, consider the following two cases from Schoenfield (2016, p. 262f.), which I will refer to throughout the thesis:

**Cheryl:** Cheryl is pregnant. She and her partner suddenly realize that, if the pregnancy is carried to term, they’ll have to skip a much-anticipated vacation that they had long ago planned with a group of friends. They don’t think that skipping the vacation would have a significantly negative impact on their lives. But, all things considered, they’d prefer that the fetus not be born. It is permissible to abort after one day for these reasons. It is not permissible to abort after nine months for these reasons. Is it permissible to abort at 150 days? 151? 151.5? Plausibly, we can create a Sorites series, admitting of borderline cases of permissibility, out of a series of abortions in which the fetus’s age differs by a day (or a minute, or a second).

**Darryl:** Darryl is watching his two-year-old daughter play in a city park. It is permissible to divert his attention from her for one second. It is not permissible to divert his attention from her for five minutes. Is it permissible to divert his attention for 30 seconds? 31? 32? Plausibly, we can create a Sorites series, admitting of borderline cases of permissibility, out of a series of diversions whose lengths differ by a second.

Here are a few more cases from the literature on moral vagueness:

**Amputations:** It is impermissible to amputate a person’s arm to save another’s life. It is permissible to amputate a person’s arm to save a billion lives. How many lives must be at stake for it to be permissible to amputate someone’s arm? Plausibly, we

can create a Sorites series, admitting of borderline cases of permissibility, out of a series of amputations, each of which is performed to save an increasing number of lives. (Schoenfield, 2016, p. 263).

Imagine there is a community jar with 5,000 pennies in it. It is not wrong (i.e. permissible) to take a few pennies from the jar when one needs some change. From this simple convention, we can generate the following sorites: P1) Taking one penny is not wrong. P2) If taking one penny is not wrong, then taking two pennies is not wrong. P3) If taking two pennies is not wrong, then taking three pennies is not wrong. ... C) Taking 5,000 pennies is not wrong. (Baima, 2014, p. 595)

There seems no precise number of milliseconds of pleasure that would compensate someone for a broken pinky toe; nor does there seem to be an exact maximal number of cents that you must spend on a taxi in order to keep your promise to meet someone at a restaurant at 7 pm. (Dougherty, 2014, p. 353)

Suppose moral responsibility consists in a cluster of rationality, self-governance and a capacity to act on the basis of moral considerations. One might hold that paradigmatically morally responsible agents have all features, but borderline morally responsible agents only have two of these features. If so, the term “morally responsible agent” would be vague. (Dougherty, 2014, p. 353)

A person's generosity . . . is a matter of degree. If the frequency of her contributions diminish, and their amount dwindles, and the grace with which she gives continually fades, then at some point she clearly is no longer generous. Somewhere in the transition, however, we might be irremediably puzzled about whether to count her a generous person. (Shafer-Landau, 1994, p. 85)

Consider a much-discussed issue in the ethics of abortion: whether (or under what circumstances) a mother is obligated to bring the fetus to term, provided the fetus is a person. One criterion that determines maternal obligation in such cases is whether the mother is responsible for creating the situation of fetal dependence. Surely a mother who had intercourse, all the while hoping to conceive, would bear the requisite responsibility. Just as surely, a victim of rape would not. But what of those

who employ a contraceptive method that is only ninety-five percent effective? Seventy-five? Forty-five? When a woman is borderline responsible for her pregnancy, she may also be borderline obligated to carry it to term. (Shafer-Landau, 1994, p. 88)

Many philosophers believe that cases like these illustrate that our moral vocabulary is vague. And once one starts looking for it, moral vagueness seems to be everywhere, including in some of our hardest decisions. Which career should I choose? How much money should I donate (or save)? When does love become excessive? Where exactly is the border between being not honest enough and being appropriately honest, and between being appropriately honest and being too honest? Should I rescue my best friend or two (three, four, five, etc.) strangers? To what extent must a nation facilitate migration by opening its borders? What is the exact number of days of imprisonment appropriate to punish a criminal? At first sight, all these problems seem to indicate that we should take moral vagueness seriously.

This thesis investigates whether that's right. I will argue that we, indeed, must take moral vagueness seriously. This verdict is—maybe contrary to how it appears at first—substantial and important.

Interestingly enough, discussions about moral vagueness often merely advocate conditional claims. As an example, consider Schoenfield's influential (2016). There, Schoenfield argues that “*if* moral vagueness exists, and robust moral realism is true, the vagueness is ontic” (p. 262, my emphasis). Wisely, I'd say, Schoenfield does not put forward an argument for the view *that* moral vagueness exists.

Couldn't one insist that cases like Cheryl's and Darryl's are sufficient to demonstrate moral vagueness? No. While many philosophers believe that these cases illustrate that the predicate ‘... is morally permissible’ is vague, not everyone is convinced.

Many would concede that in light of cases like Cheryl's and Darryl's, the predicate ‘... is morally permissible’ *seems* vague. Appearances, though, can be deceiving. Some might then simply shrug off the idea that there's moral vagueness. Moral vagueness may be theoretically interesting, but in the end, it is merely a speculative possibility. Or so those people might say. After all, some would react to Cheryl's case by saying either that she clearly is or is not permitted to abort. Probably even more would respond by saying that we just do not (yet) know whether she's permitted to abort but that we *can*, in principle, find out when ‘... is morally permissible’ applies. And they might say the same

for other alleged illustrations of moral vagueness. We just need more conceptual clarification (or maybe empirical research)! Contrast this reaction to reactions toward cases that are generally agreed upon as illustrations of vagueness. Imagine a person (let's call him Herbert) in the grey area between being bald and being hirsute. Most philosophers would agree that the question of whether Herbert is bald is, in some sense, *inquiry-resistant*. As Sorensen (2023) expresses it, a statement such as 'Herbert is bald' is "a statement that resists all attempts to settle whether it is true or false." Using another standard example, Sorensen writes: "No amount of conceptual analysis or empirical inquiry can establish the minimum number of grains needed for a heap." So far, the case has not been made that the question of whether, for instance, Cheryl is permitted to abort must also be treated as a question that resists inquiry. And many would initially disagree that it is.

Indeed, the problem for those who want to defend that there's moral vagueness is even worse. Here's why. Not only is it not clear whether alleged cases of moral vagueness *demonstrate* moral vagueness, but there's no straightforward way to *argue* for moral vagueness.<sup>1</sup> Or so I aim to show in this thesis. This, then, is the first claim I wish to defend:

*Claim 1:* There's no straightforward way to argue for moral vagueness.

Apart from this first negative claim, I'll defend two positive claims and one further negative claim. Together, these claims amount to a rationale for the view that we must take the consideration that there's moral vagueness seriously and cannot brush it off as a merely interesting theoretical possibility. Claims 2 and 3 are these:

*Claim 2:* Because it seems that moral predicates are vague, we have reason to believe that moral predicates are vague.

*Claim 3:* Vagueness has a distinct phenomenological profile.

Some might agree with Claim 2 but nevertheless not take the consideration that there is moral vagueness seriously. For several reasons, those people might insist that the appearance of moral vagueness might be deceiving. One attempt to convince them is to

---

<sup>1</sup> I will explain in chapter 3 what I mean by 'straightforward' arguments for moral vagueness.

point out that moral predicates give rise to the same phenomenology as vague predicates (Claim 3). But many people would still bring forth reasons to dismiss the existence of moral vagueness. Others might shrug off moral vagueness by claiming that moral vagueness is not especially interesting for philosophical theory and can be handled easily. Thus, I'll need to defend a fourth claim:

*Claim 4:* Apparent reasons to dismiss moral vagueness despite the appearance that moral predicates are vague are unconvincing.

Together, these four claims establish that we must take moral vagueness seriously: we have good reasons to believe that there's moral vagueness, and moral vagueness is not easy for philosophers to handle.

Establishing that moral vagueness must be addressed by philosophers, especially ethicists, is important for a number of reasons. Most crucially, many metaethical theories in standard formulations might not be able to adequately treat moral vagueness. There's an ongoing discussion on how stance-independent moral realism could account for moral vagueness (see, for instance, Schoenfield, 2016; Sud, 2019; Hawthorne, 2022). Baima (2014) argues that metaethical non-cognitivists are in trouble if moral predicates are vague. Dougherty (2014) raises problems for stance-dependent realists. Kultgen (2018, p. 53) points out that any theory that identifies moral terms with precise natural terms is at odds with moral vagueness. Moreover, addressing moral vagueness is of enormous practical significance. We need practical guidance on what to do when it's vague what to do. I'll not address these questions in this thesis, however. Instead, my aim is (more modestly) to show that they *should* be addressed.

This thesis is structured as follows. Chapter 2 clears up a common initial confusion that arises from the ambiguity of the notion 'moral vagueness'. I'll briefly discuss two ways one might understand 'moral vagueness' that I will *not* focus on in this thesis: the ethical implications of vague communication (2.1) and ethical indeterminacy arising from sources other than vague predicates (2.2).

Chapter 3 discusses why straightforward arguments for moral vagueness seem deficient (Claim 1). Straightforward arguments for moral vagueness start from the premise that moral predicates have certain features deemed sufficient for vagueness. In section 3.1, I lay out the *characterization problem of vagueness*: the problem of saying what vagueness consists in. Because of this problem, 'straightforward' arguments for

moral vagueness are not straightforward at all. Section 3.2 details why particular instances of straightforward arguments are not convincing.

Chapter 4 is the main dish. While we cannot straightforwardly argue that moral predicates are vague, we can show that moral vagueness is a philosophical datum. The central idea of this chapter is this: In the face of the characterization problem, we should appeal to the seeming that moral predicates are vague to argue that we must take moral vagueness seriously. And this seeming is enough to have a good reason to believe that moral predicates are vague (Claim 2). Sections 4.1-4.3 clarify concepts necessary to understand the methodological framework of Bengson, Cuneo, and Shafer-Landau (2022, hereafter B, C & S-L), which I will use throughout this thesis. Section 4.1 explains what, according to B, C & S-L, a philosophical datum is. Here, I also explain why I chose to use the vocabulary from B, C & S-L's theoretical framework (for instance, 'philosophical data') and how far my arguments depend on the correctness of that framework. Section 4.2 clarifies what is meant by the type of mental state called 'seeming.' Section 4.3 argues that we have specific linguistic seemings and, in particular, seemings of vagueness. Crucially, we can have seemings of vagueness without fully understanding the nature of vagueness. Section 4.4 reflects on the phenomenology of vagueness (Claim 3), providing further support for the view that it seems that there's moral vagueness.

Chapter 5 discusses how moral theories could handle the datum that moral predicates are vague. This discussion constitutes an argument for Claim 4: moral vagueness is hard to handle. Section 5.1 shows the problems that moral theorists face when trying to pass on the datum to philosophers of vagueness. Section 5.2 discusses whether Kultgen's (2018) 'Vagueness Last' approach provides what B, C & S-L call a *data-disabler*. More specifically, I argue that Kultgen does not provide strong reasons to believe that the way the datum that there's moral vagueness was produced was faulty. Section 5.3 discusses a different data-disabler: Namely, Chang's (2001) view that borderline cases can be stipulated away. As I'll discuss, if strong enough, this view would support the position that there's no moral vagueness. But I believe it isn't.

Chapter 6 is a concluding summary.

## Chapter 2: VAGUENESS IN ETHICS

This thesis investigates the question of whether we must take moral vagueness seriously. What does this question ask?

Let's think the question through by examining the predicate '... is vague'. What could we slot in this predicate's gap to catch an ethicist's attention? Three options come to mind.

Firstly, we could use *terms that refer to uttered or written expressions*. For example:

Your answer 'I met someone' is vague.

Scholz's speech is vague.

Texting 'Let's meet another time' is vague.

When we use '... is vague' like this, we want to say that a particular expression or communication style is *imprecise*. Imprecise communication has to do with morality in several interesting ways, which I will briefly address in section 2.1. Fascinating though this topic is, the question of the moral implications of imprecise communication is not the topic of this thesis.

Secondly, we could slot in *terms that refer to moral propositions, judgments, or sentences*. For instance:

The moral proposition <My future great-great-grandson has the right to live in a sustainable environment> is vague.

The moral judgment 'Alex should give Blake a compliment' is vague.

The moral sentence 'Donating to charity is morally better than helping your children with their minor problems' is vague.

We want to attribute *indeterminacy* when we apply the predicate '...is vague' to terms that refer to a proposition, judgment, or sentence. We then *broadly* use the predicate '... is vague' as equivalent to '... is indeterminate'. The focus of this thesis will not be moral indeterminacy. Instead, it will focus on whether there is vagueness in a *narrow sense* (More about this in a second).<sup>2</sup> However, vagueness in the narrow sense is possibly a source of indeterminacy. So, the main question of this thesis is strongly related to the question of whether there is moral indeterminacy. Section 2.2 will shortly lay out some potential sources of moral indeterminacy other than vagueness in the narrow sense.

---

<sup>2</sup> I borrow the notions of broad vagueness and narrow vagueness from Williams (2012).

Thirdly, we could apply ‘... is vague’ to *terms that refer to moral predicates*. For instance:

The predicate ‘... is morally permissible’ is vague.

The predicate ‘... is a morally responsible person’ is vague.

‘... is morally wrong’ is vague.

*Predicate vagueness* is vagueness in the narrow sense, and moral predicate vagueness is the focus of this thesis.<sup>3</sup> After this chapter, I will, for brevity’s sake, use the notion ‘Moral vagueness’ to refer exclusively to moral predicate vagueness.

But let me now briefly discuss the two ways one might understand ‘moral vagueness’ that I will *not* focus on in this thesis.

## 2.1 *Vagueness in Communication*

Vagueness in communication is an underexplored but important issue in ethics. Imagine that, unannounced, your partner comes home late at night. When you ask where they’ve been, the reply is, “Out.” When you further inquire what they did, the answer is, “I met someone.” When your partner gives such underspecified answers, even if they are entirely true, you might still complain that the provided information (or your partner’s communication style) is vague.<sup>4</sup>

Vague language, or *informational vagueness*, has a bad reputation. Often, when we claim that someone’s talking vaguely (for instance, by providing only general information or being ambiguous), we mean to say that that person is doing something wrong.

There are at least three ways vague language can be objectionable. First, talking vaguely can be a form of deception. Suppose in the case above that your partner has come home after having an affair and doesn’t want you to find out about this; they want to

---

<sup>3</sup> Some believe that besides predicate-vagueness, there’s also quantifier-, adverb-, and object-vagueness (for instance, Weatherson, 2010). The term ‘many’, for example, is often construed to be vague. However, ‘many’ might also be construed as a qualifier that, when combined with a predicate, forms a new, vague predicate, for example ‘... is many miles away’. The question of whether all alleged examples of non-predicative vagueness can be reduced to predicate vagueness in such a manner impresses me as very difficult, and I cannot further discuss it here.

<sup>4</sup> See Burns (1994).

deceive you without speaking falsely. Another way to put the point is that your partner obeys the Gricean quality supermaxim (“Try to make your contribution one that is true”) while violating the first quantity maxim (“Make your contribution as informative as is required for the current purposes of the exchange” (Grice, 1989, pp. 26f.)).

Some authors believe that deception is a form of lying.<sup>5</sup> However, the standard view is that a lie is an assertion the speaker believes to be false. On that view, vague language may often be an attempt to deceive without lying—as a false implicature, an attempt to mislead, or a palter.<sup>6</sup>

Second, vague language can be objectionable, not in virtue of what it conceals, but in virtue of what (and how) it reveals. Imagine your partner in the above case comes home late after an affair and *wants* you to know this. Then the answers “Out!” and “I met someone” aren’t meant to deceive. Your partner flouts the Gricean quantity maxim to communicate that they had an affair. This version of the case resembles an example provided by Grice himself: a professor writes a letter of recommendation for student X, saying, “Mr. X’s command of English is excellent, and his attendance at tutorials has been regular” (1989, p. 33). Such conversational implicature can be hurtful language. Your partner might want to hurt your feelings, and the professor might want to ridicule the student.

Third, even when vague speech isn’t intended to conceal or reveal information hurtfully, it can still be objectionable. One might answer “Out” to the question “Where have you been?” out of laziness or inattention. While such indifference must not be a form of deception or hurtful language, it might still be harmful.

So, vague language can be objectionable in at least three ways. But vague language isn’t *always* objectionable. Vague language can be more appropriate than precise language. As Brown and Levinson (1987, Ch. 5.5.2) observe, talking vaguely is often more polite than talking precisely.

Indeed, Grice’s quantity maxim can be violated not only by giving too little but also by giving too much information. Jucker et al. (2003, p. 1750) believe that one of the social functions of vague language is to “maintain an informal atmosphere in order to establish more camaraderie.” Cucchi (2007) holds that vague language displays iconicity—i.e., one can talk vaguely to suggest that more could be said. And certainly, doing so can be

---

<sup>5</sup> See, for instance, Ekman (1992).

<sup>6</sup> See Schauer and Zeckhauser (2009).

appropriate in some communicative situations. Koester (2007) indicates that vague language can signal flexibility. If your boss proposes, “Let’s meet weekly or something like that from here on out,” she might intentionally leave the frequency of future meetings open.

Furthermore, informational vagueness can be appropriate when it’s appropriate to avoid commitment: one can hedge (i.e., soften one’s words) through vague language.<sup>7</sup> Moreover, vague language sometimes is—somewhat paradoxically—more informative than precise language. Let me adopt an example from van Deemter (2010, p. 264). You know that there’s a 1.93721m terrorist at the airport, that all the other people there are smaller, and you want to inform the police. Surely, ‘Arrest the tall man’ will be more effective than ‘Arrest the 1.93721m man’.

Finally, vague language seems appropriate to report predicate vagueness, indeterminacy, and ignorance. Vague statements like ‘It’s sort of permitted to abort at 151 days’, ‘Jones is kind of a morally responsible person,’ and ‘Just be kind!’ might be precisely the right thing to say in some cases.

The preceding discussion suggests that vague communication might sometimes be objectionable and sometimes unobjectionable or even morally required. But many unanswered questions remain. Can vague language be wrong independently of its negative consequences? Many philosophers believe that deception can constitute a violation of autonomy. Can similar accounts of the wrongness of vague speech be developed? Moreover, one wonders if there are general principles governing the ethics of vague communication. Such questions might be fruitfully explored in future research, but I will not further pursue them here.

## 2.2 *Moral Indeterminacy*

Moral vagueness, in the narrow sense, is moral predicate vagueness. However, some authors use ‘vagueness’ broadly to refer to indeterminacy.<sup>8</sup> Does ethics have pockets of indeterminacy?<sup>9</sup> I now present three possible sources of indeterminacy apart from predicate vagueness that plausibly occur in ethics. This is not the place to study how these

---

<sup>7</sup> See Milanović and Milanović (2010).

<sup>8</sup> See, for instance, Rabinowicz (2009).

<sup>9</sup> Pravato claims (overconfidently, I’d say) that indeterminacy is “a totally unremarkable fact about normative language” (2022, p. 2139).

different forms of indeterminacy relate to each other. Could some forms of moral indeterminacy be reduced to others? Maybe! Unfortunately, I have no space to pursue this fundamental matter further.

First, consider contingent propositions about the future. Łukasiewicz famously voiced a view that is not uncommon today when he claimed that, if the future is undetermined, “‘I shall be in Warsaw at noon on 21 December of next year’, can at the present time be neither true nor false” (2005, p. 53).

The truth value of contingent propositions about the future is, for instance, relevant to the debate about the rights of future persons. <My grandson shall have the right to live in a healthy environment> might—just as <I shall be in Warsaw at noon on 21 December of next year>—be neither true nor false. Macklin probably has something like this in mind when discussing the view that “the ascription of rights is properly to be made to actual persons—not possible persons” (1981, p. 151). That’s one way to arrive at the suspicion that there’s indeterminacy in ethics.

Next, consider this case: Alex wants to be kind to Blake this evening. So, Alex asks himself: ‘Should I give Blake a compliment? Or offer him a drink? Or a snack?’ Some moral duties seem overly general. The duty of beneficence, which Alex wants to fulfill, doesn’t precisely tell him what to do. And—as Dougherty suspects— “it may be that the global structure of one’s duties is determinate, but the local structure is not” (2014, p. 352). <Alex should give Blake a compliment> might thus be another example of an indeterminate moral proposition.

Now, consider a classic illustration of incomparability given by Raz (2009, p. 332). You face two options: being equally suited for both and having an equal chance of success in both, one commits you to a career in law, the other to a career as a clarinetist. It doesn’t seem that one career path is better than the other. Yet there’s intuitive resistance to the idea that they are equally good. When facing such a choice, Raz claims, “reason is indeterminate. It provides no better case for one alternative than for the other” (pp. 332f.). But, I hasten to add Raz here doesn’t want to say that the statement ‘It is better to become a clarinetist than to become a lawyer’ is indeterminate. If incomparability obtains, “it is false that of A and B either one is better than the other or they are of equal value”; if indeterminacy obtains, “this statement is neither true nor false” (p. 324). For Raz, then, incomparability doesn’t entail indeterminacy.

However, Broome offers an influential attack on Raz’s position. Imagine a series of variations of a legal career, ordered by how much better they are compared to a musical

career. Broome argues that incomparability entails that in such a series, there will be “points such that it is neither true nor false that they are *F*er than the standard [better than the musical career], and neither true nor false that the standard is *F*er than them” (1997, p. 131). Broome calls the zone of such points a zone of ‘soft indeterminacy’ or vagueness. If Broome is correct, the indeterminacy of moral propositions or sentences that include the better than relation is thus another pocket of moral indeterminacy (or moral vagueness in a broad sense).<sup>10</sup>

If there are different forms of moral indeterminacy, there’s a need to systematize them. Some guiding questions for such a project might be: Are some indeterminacies reducible to other forms? How can one clearly distinguish between different forms of moral indeterminacy? I suspect that different forms of moral indeterminacy might call for different responses on the part of agents. For instance, different forms of moral indeterminacy might come with different second-order “contrary-to-determinacy imperatives” (Sorensen, 1991, p. 306). Williams (2016, 2017) offers a detailed analysis of what to do when it’s indeterminate what to do, but a thorough differential treatment of kinds of moral indeterminacy is still outstanding. These matters need a careful, separate investigation—something I cannot embark on here.

### Chapter 3: **STRAIGHTFORWARD ARGUMENTS FOR MORAL VAGUENESS**

From now on, let me use the notion ‘Moral Vagueness’ to refer to moral predicate vagueness for brevity’s sake. That said, the question of whether there’s moral vagueness then boils down to this question: Are moral predicates like ‘... is morally permissible’, ‘... is morally forbidden’, ‘... is virtuous’, ‘... is a morally responsible person’, and so on, vague? At first, it might seem that one can give a straightforward answer. Yes, moral predicates are vague because they have features sufficient for being vague. Here is an

---

<sup>10</sup> Apart from the intuition that predicate vagueness, future contingencies, generality, incomparability, or some other source of indeterminacy gives rise to moral indeterminacy, stance-independent moral realists have a further reason to accept moral indeterminacy. Shafer-Landau argues that embracing indeterminacy in ethics allows moral realists to explain the appearance of deep moral disagreement. Even ideal reasoners, Shafer-Landau suggests, “can fail to come to a consensus on which option is uniquely best, because, owing to indeterminacy, there may be no such option” (1994, p. 343).

---

argument template that presents a straightforward way to argue that moral predicates are vague.

**P1** If a predicate has feature *a*, it is vague.

**P2** Moral predicates have feature *a*.

Thus, moral predicates are vague.

In this chapter, I explain why arguing for moral vagueness in this way is not straightforward at all. My basic idea is simple: Straightforward arguments presuppose that there's theoretical consensus of what vagueness consists in. That's, however, not the case: we don't know which features make a predicate vague. I will first give a short overview of the so-called characterization problem of vagueness. I will then argue why particular instances of the straightforward argument fail.

### **3.1** *The Characterization Problem of Vagueness*

The characterization problem of vagueness is the problem of saying what the vagueness of predicates consists in. Some authors (Greenough, 2003; Bueno and Colyvan, 2012) believe that finding a mutually agreed-upon definition of vagueness is needed to avoid crosstalk between proponents of different theories of vagueness. Others wonder whether a theory-neutral characterization of vagueness is possible or desirable (Smith, 2005; Weatherson, 2010). In any case, it is surprisingly hard to characterize the nature of vagueness. The three central attempts are to characterize vagueness via the feature of borderlineness, via the feature of boundarylessness, or via the feature of giving rise to specific epistemic or psychological responses. The following subsections will focus on the points of disagreement between those attempts.

#### **3.1.1** *Vagueness as Borderlineness*

The first idea is that vague terms are characterized by permitting borderline cases. Borderline cases are traditionally construed as cases of indeterminacy—that is, as cases to which a predicate neither truly nor falsely applies, resulting in a truth-value gap. Fine (1975), for instance, illustrates vagueness through the (artificial) predicate ‘... is nice<sub>1</sub>’. Per stipulation, a number *n* is nice<sub>1</sub> if  $n > 15$ , and not nice<sub>1</sub> if  $n < 13$ . Therefore, ‘14 is nice<sub>1</sub>’ lacks a definite truth value, and the number 14 is a borderline case of ‘... is nice<sub>1</sub>’.

Alternatively, some authors (for instance, Tye, 1994) hold that sentences with vague predicates can have one of three truth values: true, false, or indeterminate.

There are two immediate concerns with this approach. First, if borderliness is the feature of a term to have cases that result in truth-value gaps on predication, the present king of France seems to be a borderline case of ‘... is bald’, and that indicates that ‘... is bald’ is vague. But it’s commonly assumed that if ‘The present king of France is bald’ lacks truth-value, this is due to presupposition-failure, not due to the vagueness of ‘... is bald’. To argue that ‘... is bald’ is vague, one would have to appeal to *special* borderline cases (such as Harry, the borderline bald man). In general, *mere* borderliness seems to be an inadequate distinguishing feature between different sources of indeterminacy (Schiffer, 2000; Eklund, 2005).<sup>11</sup>

Second, characterizing vagueness in terms of indeterminacy is not theory-neutral. At least epistemicists like Williamson (1996)—who reject the idea of vagueness-related truth-value gaps—would protest. Weatherson (2010) holds that his characterization of vagueness as indeterminacy is compatible with an epistemic notion of indeterminacy but does not expand on this point. I will further explain epistemicism and other theories of vagueness in section 5.1. For now, I’ll move on.

### 3.1.2 *Vagueness as Boundarylessness*

On to the second attempt. Sainsbury (1996) influentially expressed the intuition that predicates like ‘... is nice<sub>1</sub>’ are not vague because they have precise boundaries. However, characterizing vagueness via boundarylessness has a critical problem—to many, boundarylessness seems a confused concept. Imagine hair after hair growing on a bald person’s head. It seems unintelligible that there’s no boundary at which that person turns non-bald, and appealing to a vague boundary instead would be circular (Eklund, 2005). That’s a steep challenge for the attempt to characterize vagueness via boundarylessness or related concepts.

Boundarylessness is closely related to tolerance and sorites-susceptibility (or soriticality). A predicate is tolerant iff small changes in a relevant underlying property are not sufficient for making a change to whether it applies to a case (Wright, 1975), and

---

<sup>11</sup> A further complication is that indeterminacy has its own characterization problem (Taylor, 2018). So, characterizing vagueness in terms related to indeterminacy might be to jump out of the frying pan into the fire.

sorites-susceptible (or soritical) iff it gives rise to the sorites paradox (Bueno and Colyvan, 2012). There are many ways to formulate the sorites paradox. I prefer the following one:

**P1** Predicate  $F$  has an underlying property relevant for its application, and there's a case of  $F$ -ness with a certain value of that underlying property.

**P2** There's a case of not- $F$ -ness with a different value of that property.

**P3** Small changes in the underlying property are not sufficient for making a change to whether  $F$  applies to a case [Tolerance Principle]

A predicate  $F$  gives rise to a sorites paradox if these three claims appear plausible. These three claims, however, also imply (absurdly, most would say) that there are cases that are  $F$  and not- $F$ .

To illustrate:

**P1** The property of having a number of cranial hairs is relevant for the application of '... is bald'. People who have only one hair on their head are bald.

**P2** People who have 20,000 hairs on their heads are not bald.

**P3** One hair more or less cannot make a difference in whether someone is bald. [Tolerance Principle]

Therefore (absurdly): People with only one hair on their heads are bald and not-bald; people with 20,000 hairs on their heads are not-bald and bald.

Characterizing vagueness via tolerance or sorites-susceptibility is prone to produce inconsistencies. The Sorites is a paradox, after all. Eklund's (2005) answer to this problem is that the meanings of vague terms are, in fact, inconsistent: vague terms are characterized by the disposition of competent speakers to accept tolerance, which leads to the above-illustrated absurdities.

In contrast, Smith's (2005) characterization via closeness is meant to capture the intuitions that might lead one to characterize vagueness via boundarylessness/tolerance/sorites-susceptibility without falling into inconsistencies. Roughly, a predicate exhibits closeness iff small changes in a relevant underlying property correspond to small changes in respect to truth. However, not everyone is on board with the idea that there are degrees of truth. Moreover, Smith's view fails to account

for what Weatherson calls “lumpy” vagueness: it seems that for some vague predicates (for instance, ‘... is late’), small changes in a relevant underlying property can correspond to sudden changes in respect of truth (Weatherson, 2004).<sup>12</sup>

### 3.1.3 *Epistemic and Psychological Characterizations*

Finally, one might attempt to characterize vagueness via epistemic/psychological notions. My phenomenological description of vagueness in section 4.4 will be partially informed by some of these attempts. Note that ‘borderline case’ in the following does *not* refer to cases of indeterminacy. Instead, ‘borderline case’ here refers to cases that, when we apply the corresponding predicate, result in a proposition towards which we have specific mental attitudes.

Since not all (even inescapable) *ignorance* is due to vagueness, one cannot sufficiently characterize vagueness via an unqualified epistemic notion: borderline cases are not simply cases for which it’s (even inescapably) unclear whether a term applies to them (Peirce, 1902).

Instead, one might attempt to flesh out the intuitive idea that vagueness gives rise to *ambivalence*. Schiffer (2000) proposes that borderline predications are such that one can have a *vagueness-related partial belief* (VPB) about them. Unlike standard partial beliefs (or credences), VPBs do not conform to the laws of classical probability. For instance, when *S* has a 0.4 VPB that *p* and a 0.6 VPB that *q*, it does not follow that *S* should have a 0.24 but a 0.4 VPB that *p* & *q*. Among others, Edgington (2016) has pointed out some unintuitive results of such VPB rules.

Wright’s (2001) way to flesh out the ambivalence-intuition is that borderline predications give rise to *quandary*: “We do not know whether to endorse them, do not know how we might find out, and can produce no reason for thinking that there is, or even could be, a way of finding out [whether they are true or false]” (p. 79). However, Wright concedes that not only borderline predications can give rise to quandary (think, for instance, of Goldbach’s conjecture).

Greenough (2003) proposes to characterize vagueness instead in terms of *epistemic tolerance*. The idea is that a predicate is vague iff small changes in a relevant dimension

---

<sup>12</sup> If there is indeed ‘lumpy’ vagueness, this might defuse the hypersensitivity objection to epistemicist moral realism (see section 5.1).

are insufficient for making a known difference as to whether a predicate applies to a specific case.<sup>13</sup>

An objection to all attempts to characterize vagueness via epistemic/psychological notions is that such attempts merely capture an *effect*, not an essential characteristic (Eklund, 2005). What to think about this objection depends on the goals of characterizing vagueness. Do we want to capture the essence of vagueness? In that case, it seems indeed true that epistemic/psychological characterizations are not enough.

However, the upshot of the characterization problem is exactly that we do not know the essence of vagueness. Maybe we never will. But there's another way to understand 'characterizing vagueness.' We might understand 'characterizing' in a phenomenological sense. The key idea is simple. When we do not know what something essentially is, we can still describe what it's like to experience it. Consider a parallel. I don't know what beauty essentially consists in. But I can describe what experiencing beauty is like. I can, for instance, describe the mental states (for instance, admiration, awe, attraction) we typically have towards instances of beauty. And this goes a long way toward distinguishing beauty from other properties. So, for the task of phenomenological characterization, it's apt to describe the characteristic mental states we have towards borderline predications.

I will turn to the phenomenology of vagueness in section 4.4. For now, let me address in more detail particular straightforward arguments.

### 3.2 *Particular Straightforward Arguments*

When philosophers argue for the existence of moral vagueness, they typically point out that moral predicates display some of the following features:<sup>14</sup>

*Imprecise gradability:* A vague predicate admits degrees, but one cannot precisely measure those degrees.

---

<sup>13</sup> Furthermore, Weatherson (2010) criticizes both Smith and Greenough for not taking non-predicative vagueness into account.

<sup>14</sup> The following list has the same members that Constantinescu calls (wisely) not features but "standard symptoms of vagueness" (2014, p. 157f.).

*Boundarylessness*: There's no sharp boundary dividing cases to which a vague predicate applies from cases to which it doesn't.

*Borderlineness*: A vague predicate admits of borderline cases.

*Tolerance*: Small changes in relevant properties do not change whether a vague predicate applies to a case.

*Soriticality*: A vague predicate gives rise to the sorites paradox.

It should be clear from the preceding section that none of these alleged features is uncontroversially sufficient for vagueness. To repeat, there's no consensus over what vagueness consists in. Those who characterize vagueness via boundarylessness will usually reject borderlineness as a sufficient condition. And many of those who characterize vagueness via borderlineness would reject boundarylessness and its related notions as incoherent. Due to this deep lack of consensus, all straightforward arguments fail as we lack justification for endorsing **P1** (If a predicate has feature *a*, it is vague).

Yet it's worth showing in more detail why particular 'straightforward' arguments are not straightforward at all.

### 3.2.1 *The Straightforward Argument from Imprecise Gradability*

To my knowledge, no standard theory of vagueness takes imprecise gradability alone as a sufficient condition for vagueness. Nevertheless, let us start with the straightforward argument from imprecise gradability.

**P1** If a predicate is imprecisely gradable, it is vague.

**P2** Moral predicates are imprecisely gradable.

Thus, moral predicates are vague.

First, consider Dougherty's (2014) example of Jones, who has some but not all of the cluster of capacities that morally responsible persons have. According to Dougherty, someone like Jones might be considered a borderline case of moral responsibility. Dougherty also seems to believe there's a difference between cluster predicates such as '... is morally responsible' and gradable predicates. Moreover, Dougherty apparently believes that borderlineness is sufficient for vagueness. He thus concludes that a case like Jones's shows that gradability is unnecessary for explaining vagueness. And if gradability

is not a necessary feature of vague predicates, *imprecise* gradability is also not a necessary feature of vague predicates. Similarly, other authors (for instance, Bueno & Colyvan, 2012) have argued that non-gradable predicates such as ‘... is a religion’ can have borderline cases (and are arguably thus vague) because there can be cases where some but not all necessary criteria for being a religion are fulfilled. However, these are reasons to reject imprecise gradability as a *necessary* feature of vague predicates. To reject **P1**, we need to argue that imprecise gradability is insufficient for vagueness.

We can do so. Some predicates are gradable and nevertheless have an on/off characteristic. Plausibly, on/off predicates should not be considered vague. For instance, ‘... is pure’ admits of degrees (this bar of gold is less pure than that) but does not seem vague (it’s clear when a bar of gold is pure). Now, ‘... is pure’ might be considered a *precisely* gradable predicate. But nothing seems to preclude the possibility that *imprecisely* gradable predicates can also be on/off and, therefore, not vague. Maybe ‘... is closed’ and ‘... is empty’ are examples of imprecisely gradable predicates, yet on/off and thus not vague.

Now, let’s return to Dougherty’s example. One could claim that Jones is (imprecisely) less morally responsible than Smith but that neither of them is *really* morally responsible. ‘... is morally responsible’ might be imprecisely gradable but on/off. Or, to avoid the contrastive use of ‘... is morally responsible,’ maybe both Jones and Smith are *somewhat* but not *really* morally responsible. The analogy to predicates such as ‘... is pure’, ‘... is closed’, and ‘... is empty’ is perhaps even more compelling for deontic predicates such as ‘... is morally permissible’ because such predicates are traditionally conceived as on/off terms. As Tappolet (2013, p. 5) observes, “Ordinary language does not allow for comparative and superlative forms of deontic concepts.” Even if one would claim (at odds with ordinary language) that it’s (imprecisely) less permissible for Cheryl to abort at 151 days of gestational age than after 150 days or that it’s *somewhat* permissible in both cases, that doesn’t necessarily mean it’s *really* permissible to abort at either time.

### 3.2.2 *The Straightforward Argument from Boundarylessness*

Here’s the straightforward argument from boundarylessness.

**P1** If a predicate lacks boundaries, it is vague.

**P2** Moral predicates lack boundaries.

Thus, moral predicates are vague.

Apart from the inconsistencies that the idea of boundarylessness seems to come with, an interesting response to the straightforward argument from boundarylessness is the following: Think again of Dougherty's example. These seem to be fair questions about his case: Is someone who lacks one capacity that clearly morally responsible agents would have slightly less morally responsible than such an agent? Or is such a person already a borderline case of '... is a morally responsible person'? If the latter, some might claim that '... is a morally responsible person' could be a predicate *with* sharp boundaries. People with the entire cluster of capacities are in its extension; people who lack all of them are in its anti-extension; people who have at least one but not all of them are in its 'penumbral' extension (they are in the 'zone of indeterminacy'). Similarly, there might be a sharp cutoff between cases where abortion is permissible and cases where it's indeterminate whether abortion is impermissible, and between this zone of indeterminacy and cases where abortion is morally prohibited. So, even if one maintains that moral predicates have borderline cases, one must not also hold that moral predicates are boundaryless. For such reasons, it is not clear why one should accept **P2**.

### 3.2.3 *The Straightforward Argument from Borderlineness*

Let's then turn to the straightforward argument from borderlineness.

**P1** If a predicate has borderline cases, it is vague.

**P2** Moral predicates have borderline cases.

Thus, moral predicates are vague.

One should be able to define what borderline cases are to bring forth such an argument. As mentioned above, borderline cases are traditionally understood as cases of indeterminacy. This is exactly how Sud, in a recent article on moral vagueness, understands them (2019, p. 685). But, as Sud notices, if all cases of indeterminacy are borderline cases, and vagueness is just one source of indeterminacy, borderlineness does not imply vagueness in the narrow sense (predicate vagueness). Moreover, Constantinescu (2014) and Kultgen (2018) argue that there are predicate-independent borderline cases. 'This cloud weighs one billion kilograms' or 'The number of heaps in

the sandbox is prime' (imagine four heaps and one borderline heap in the sandbox) seem like borderline predications. Yet '... weighs one billion kilograms' and '... is prime' are not vague predicates. Similarly, there might be borderline cases of '... is permitted to abort' in embryonal development, but Kultgen argues that this must not imply that '... is permitted to abort' is vague (2018, p. 62). Instead, '... is an act of killing a zygote' and '... is an act of killing a person' might be vague. These constitute good reasons to reject **P1**.

### 3.2.4 *The Straightforward Argument from Soriticality/Tolerance*

Because sorites paradoxes involve a tolerance principle as a premise, we can discuss straightforward arguments from soriticality and from tolerance together.

**P1** If a predicate is soritical/tolerant, it is vague.

**P2** Moral predicates are soritical/tolerant.

Thus, moral predicates are vague.

Here's a simplified variation of the sorites paradox I already gave above:

**P1** Someone with 200,000 hairs on his head is hirsute.

**P2** Someone with one hair on his head is not.

**P3** Plucking off one hair from a hirsute person doesn't make that person not hirsute.

[Tolerance Principle]

These three claims appear plausible, yet they also seem to imply (absurdly) that someone can simultaneously be both hirsute and not hirsute.

And here's an (alleged) sorites paradox arising from Cheryl's case:

**P1** Cheryl is permitted to have an on-demand abortion after one second of pregnancy.

**P2** Cheryl is prohibited from having an on-demand abortion after nine months of pregnancy.

**P3** A one-second difference in gestational age doesn't change whether an on-demand abortion is permissible. [Tolerance Principle]

Again, these three claims are plausible but apparently imply an absurdity: Namely, that there's a moment at which abortion is permissible and not permissible.

However, appearances can be deceiving. To some, the above might illustrate a slippery slope rather than a paradox. If it were permissible for Cheryl to abort after one second of pregnancy, then (because of the tolerance principle) it would (absurdly) be acceptable to abort after nine months. Others might see a slippery slope in the other direction. So, it's difficult to identify true sorites paradoxes (in ethics and elsewhere).

Similarly, Chang says that '... breaks the camel's back' seems soritical, but we know small changes make a difference here because physics tells us there's an exact breaking point for the camel's back. True soriticity requires compelling, innocent-seeming premises, including an intuitive tolerance principle. But how could we assess that a predicate is truly soritical without appealing to vagueness? Chang concludes: "We need to determine whether a predicate is vague before we can ascertain whether a sorites-seeming argument is a genuine sorites" (2001, pp. 136f.). Thus, arguing that moral predicates are vague because they are soritical/tolerant is not straightforward at all: both **P1** and **P2** of the straightforward argument from soriticity/tolerance are problematic.

Where does this leave us? Straightforward arguments cannot establish that moral predicates are vague. Such arguments presuppose that we know what vagueness consists in. But we don't. As outlined just now, different straightforward arguments also face specific problems. And I doubt whether there can be other ways of straightforwardly arguing moral predicates are vague.

#### Chapter 4: MORAL VAGUENESS AS A PHILOSOPHICAL DATUM

Now for the main course. This chapter aims to show that we can do more than shrug our shoulders regarding the question of whether there's moral vagueness. We have good reasons to believe that moral predicates are vague. I showed above that there's no consensus about what vagueness consists in. This precludes the possibility of (non-controversially) demonstrating (via straightforward arguments) that moral predicates are vague. I will now argue that we can, more modestly, show that moral vagueness is a *philosophical datum*—a consideration that philosophers must take seriously. Even though this appears to be a humble goal, it is important to show that moral vagueness must be seriously considered. To see why, remember that the main question in the recent literature

on moral vagueness is this: Which metaethical theories are compatible with moral vagueness, and which are not? Now, imagine someone would bring forth the following argument:

**P1** There's moral vagueness.

**P2** If moral realism were true, there would be no moral vagueness.

Thus, moral realism is false.

After what I discussed in the preceding chapter, the moral realist could just shrug her shoulders. "We're not even clear about what vagueness consists in," she might respond, "so why believe **P1**?" I aim to show that the moral realist would have to do more. She must take moral vagueness seriously.

The central idea of this chapter is this: In the face of the characterization problem, to argue that we must take moral vagueness seriously, we must appeal to the *seeming* that moral predicates are vague.

Sections 4.1-4.3 clarify concepts necessary to understand this claim in light of the methodological framework of B, C & S-L (2022), which I will use throughout this thesis. Section 4.1 explains what, according to B, C & S-L, a philosophical datum is, spells out what role their framework plays in this thesis, and justifies the choice of this framework for this thesis. Section 4.2 clarifies what is meant by the type of mental state called "seeming." Section 4.3 argues that we have specific linguistic seemings and, in particular, seemings of vagueness. Crucially, we can have seemings of vagueness without fully understanding what vagueness consists in. Or so I argue. Section 4.4 reflects on the phenomenology of vagueness. The conjecture I will articulate is that vagueness has a distinct phenomenological profile. The phenomenological reflections of section 4.4 provide further support for the idea that it seems that there's moral vagueness.

#### **4.1 Philosophical Data**

In this section, I first (4.1.1-4.1.3) reproduce the view on the nature of philosophical data from B, C & S-L (2022). I will contrast how B, C & S-L answer the following questions with other possible answers: What is the nature of data? What makes something a datum?

What is the relationship between data and theory?<sup>15</sup> I will then (4.1.4) clarify and justify my appeal to B, C & S-L's framework.

#### 4.1.1 *The Nature of Data*

Let's first consider the nature of data—or, if you will, the ontological status of data.

To illustrate, imagine a Sherlock Holmes case where Holmes is in the dark about who the murderer is. He examines a mysterious letter through his magnifying glass and spots a peculiar black dot. He notes, 'Discovered a strange black dot on the letter!' He collected the first datum for his theory about the case. But what exactly is the datum? One view is that the datum is his *finding* (or observation) of a black dot on the letter. Another view is that the datum is the *record* of his finding: the entry in his notebook.

Here's another illustration. Imagine a scientist examining an unknown material with the help of a thermometer. She takes a photo of the thermometer when the material turns liquid. According to the first view, the datum is the observation that the material melted at 327.4 °C. According to the second, it's the photo. These two views differ drastically in what ontological status they ascribe to data. According to the first view, data are (or are intimately related to) psychological states. According to the second, data are material objects.<sup>16</sup>

Now imagine that, after gathering more and more clues, Holmes deduces that the victim cannot have written the letter. Or that after enough measurements, our scientist concludes that the material melts at 327.5 °C. A third view on the nature of data is that a

---

<sup>15</sup> Let me note upfront that the term 'data' has no uniform usage. Like the term 'evidence', it is used both by philosophers and non-philosophers in a number of ways. Some of these are complementary and others are not. To add to the confusion, 'data' and 'evidence' are often used synonymously. On the other hand, many insist that these notions should be kept apart strictly. This section's main aim is to explain the way the notion of 'data' is used by Bengson, Cuneo, and Shafer-Landau (2022). I'm addressing the question of whether there's a *correct* way of using the notion of 'data' in 4.1.4.

<sup>16</sup> When thinking about the difference between findings and records of findings, the devil is in the details. Woodward (2010) claims that data are the individual results of measurement or detection procedures (possibly with the help of instruments) or, "by extension," records or reports of such results, such as photos, graphs, tables, or displays. For Hacking (1992), data are *marks*, the direct outputs of our measurement procedures. Those marks need to be *processed* (i.e., assessed, reduced, and analyzed) before they can be *interpreted*. In contrast, Rheinberger (2011) holds that data are not the direct outputs of our measurement procedures. For him, such mere marks (or *traces* in his words) need to be processed into *symbols* to count as data.

datum is such a *serious consideration*. The severe suspicion that the victim didn't write the letter is Holmes's datum for his theory about the case; the assumption that the material melts at 327.5 °C is our scientist's datum for a theory about the unknown material.

This last view on the nature of data is the view of B, C & S-L. According to these authors, data are considerations a theory must handle—that is, data are *propositions*.<sup>17</sup> More specifically: propositions relevant to a research domain (such as ethics, philosophy of vagueness, etc.) that are *fallible yet plausible*. The latter feature distinguishes data from one crucial way the notion of 'evidence' is used. As B, C & S-L point out (2022, p. 37), philosophers like Timothy Williamson maintain that evidence is factive. For Williamson, evidence is our relevant knowledge. According to B, C & S-L, a datum is, in contrast, non-factive. A datum is a relevant proposition that might, however, be false.<sup>18</sup>

#### 4.1.2 What Makes Something a Datum

Following the view of B, C & S-L on the nature of data, we can now more clearly understand what makes something a datum. The question is what makes propositions relevant to a research domain and what makes them plausible.

Some philosophers are quite down-to-earth about what makes propositions relevant for a research domain. According to them, what counts as relevant is decided by what theorists believe to be relevant to their theories.<sup>19</sup> B, C & S-L are more demanding. They believe that data is *inquiry-constraining*. A proposition must *really* relate to the subject matter to count as a datum.

Here's an illustration of this difference. Imagine all scientists would agree that a theory about electromagnetism must take into account that I drink too much coffee. According to the down-to-earth approach, it's thus a datum for theories about

---

<sup>17</sup> I'll set aside questions about the nature of propositions. I don't believe that it matters much for the theory of B, C & S-L whether propositions are abstract or mental, structured or simple, and so on.

<sup>18</sup> As already noted, the notion 'evidence' is used in numerous ways. Like the term 'data', 'evidence' is often used to refer to findings, records of findings, and interpretations of findings. Maybe it is true that 'data' is usually not used to refer to *facts*, but 'evidence' at least sometimes does. The matter of distinguishing data from evidence needs a careful separate treatment—something I cannot embark on here. However, I'd like to note that Bengson, Cuneo, and Shafer-Landau are not always clear on this distinction. For instance, they do not explain what exactly they mean when they say that data can provide evidence (2022, p. 62) or that data sets can be supported by evidence (2022, p. 66).

<sup>19</sup> A view in the vicinity of this is expressed by Leonelli (2016).

electromagnetism that I drink too much coffee. However, my coffee consumption is irrelevant to a theory about electromagnetism.<sup>20</sup> According to B, C & S-L, it's thus *not* a datum for such theories that I drink too much coffee. In other words, B, C & S-L advocate that relevance is *objective*, not constructed.

Next, plausibility. B, C & S-L maintain that if a proposition is a datum, inquirers, considered collectively, have good reason to believe this proposition (2022, pp. 56f.). B, C & S-L hold that the group of theorists must possess good reasons to believe a proposition for it to count as a datum. They intentionally leave the notion of 'good reasons' open, stating that their view "can be combined with views as divergent as virtue epistemology (which appeals to inquirers' competences), phenomenal dogmatism (which focuses on inquirers' seemings or presentations), and subjective Bayesianism (which invokes probability given inquirers' background beliefs)" (2022, p. 57).

#### 4.1.3 Data and Theory

Finally, what is the relationship between data and theory? For B, C & S-L, the answer is: data are the inputs of theorizing. They understand theorizing as an activity that aims to construct a theory that provides understanding. Their preferred phrase for what one does with data when one theorizes is that one 'handles the data.' For instance, when Sherlock Holmes theorizes about the murder case, he must somehow deal with the datum that the victim did not write the letter. More fully, B, C & S-L claim that theorizing involves *accommodating* and *explaining* the data. Alternatively, a theorist might also defend the view that specific data do not have to be accommodated or explained. That's what B, C & S-L mean when they say that theories must handle the data.<sup>21</sup>

---

<sup>20</sup> At least I hope so. I'd hate having physicists observing me at breakfast.

<sup>21</sup> It's interesting to contrast the view of Bengson, Cuneo, and Shafer-Landau on the relationship between data and theory with that of Bogen and Woodward. Bogen and Woodward (1988) influentially argued that theories do not have to explain data but *phenomena*. However, it must be stressed that they suggest calling only the results of detection processes and reports of such results data. What they call 'phenomena' is what can be inferred from such results. It seems to me that Bogen and Woodward's phenomena are Bengson, Cuneo, and Shafer-Landau's data; but that would be a substantive claim that I won't defend here. In any case, it is pretty clear that theories do not have to explain data conceived of as reports of findings. Suppose we say that the datum for Holmes' theory is the entry in his notebook. Indeed, his theory about who the murderer is doesn't have to explain all the details of his entry, for instance, why he squiggled the 'D' of

Here's what B, C & S-L mean by accommodation and explanation: "A theory T 'accommodates' a datum just in case that datum is likely to hold or be true, given T," and "a theory T 'explains'  $\phi$  just in case T invokes some  $\psi$  such that  $\phi$  holds or is true because  $\psi$  holds or is true" (2022, pp. 110, 112). Let's call accommodation and explanation the *positive way of handling data*.

The *negative way of handling data* is to defend that a specific datum does not have to be accommodated or explained. B, C & S-L mention four negative strategies for handling data.

- (i) Pointing to (at least equally strong) considerations that the way the datum was produced was faulty.
- (ii) Pointing to (at least equally strong) considerations supporting a position that contradicts the datum.
- (iii) Passing on the job of accommodating and explaining to a different theory that coheres with one's.
- (iv) Pleading unexplainability.

The considerations mentioned in (i) and (ii) are what B, C & S-L call 'data-disablers.' Here's their official definition of this notion:

A set of considerations C . . . disables a datum for a set of inquirers if and only if C provides them with good reason either to believe that the datum is mistaken or to suspend judgment with respect to it, and this reason is at least as strong as the reason(s) for believing that datum. (2022, p. 72).

Keep the ways of handling data in mind for Chapter 5. To conclude this section, let's briefly recap some key notions by returning to Holmes' murder case.

The datum Holmes is concerned with is that the victim didn't write the letter. One way a theory would *accommodate* this datum is by incorporating the idea that the letter was written after the victim was already dead. This would surely make it likely that the victim didn't write the letter! A theory can *explain* this datum by incorporating the idea

---

'Discovered'. Perhaps Bengson, Cuneo, and Shafer-Landau run into similar problems even though their ontology of data differs from that of Bogen and Woodward? Maybe! But I leave the matter there.

that the murderer wrote the letter to set Holmes on the wrong track. That would answer why it wasn't the victim who wrote the letter. But Holmes might have second thoughts. Maybe the magnifying glass he used to investigate the letter was dirty? This would be a *data-disabler of the first kind (i)*: a consideration that points out that the way the datum was produced was faulty. However, Holmes later interrogates the honest-seeming butler, who swears that he saw the victim write the letter herself. This is a *data-disabler of the second kind (ii)*: it supports the contradictory position that the victim *did* write the letter. In building his theory, maybe Holmes *passes on (iii)* some data to other theories. For instance, he might say that Freudian psychology nicely predicts that the victim's son behaves so jealously, and Holmes then leaves it to the psychologists to fill in the details. Or Holmes might brush off the datum that Moriarty tricked him again—Moriarty seems to be behind all this!—as *unexplainable (iv)*. It's just a brute fact that his nemesis is evil.

The claim I will make and defend in subsection 4.2 is that the mental state called 'seeming,' and in particular, seemings of vagueness can produce philosophical data. But before moving onto that, I should explain the role that the framework of B, C & S-L plays in this thesis and justify the choice of this framework.

#### 4.1.4 *Why Appeal to Bengson, Cuneo, and Shafer-Landau?*

A natural worry at this point is this: Why adopt the framework of B, C & S-L? More specifically, one might raise the following questions: Why endorse B, C & S-L's view on philosophical data when, as I suggested above, there are many other possible views on this concept? Why use the notion of philosophical data at all? And finally, why take the specific methodological stance they take? Why not use a more traditional method, such as Rawls' (1971) reflective equilibrium or Armstrong's (1989) intellectual cost-benefit analysis?

First, let me address the worry that the view of philosophical data put forth by B, C & S-L might be inadequate. Shouldn't I defend their view that philosophical data are objectively relevant propositions that we have good reasons to believe? I agree that this task would be important to defend B, C & S-L's theoretical stance, but it's not one I need to tackle here. Recall that this thesis aims to argue that we should take moral vagueness seriously. And B, C & S-L call considerations we must take seriously 'philosophical data.' Using their vocabulary, 'Moral vagueness is a philosophical datum' thus just restates what I aim to argue for. Yes, this restatement makes implicit claims about when

one must take considerations seriously: Namely, when they're objectively relevant propositions we have good reasons to believe. But this is a natural and appealing view. What follows in this chapter can be seen as an argument for there being good reasons to hold that moral vagueness should be taken seriously according to this view. However, I believe not much hangs on how we construe 'taking the consideration that there is moral vagueness seriously.' The argument I will bring forth provides, I think, grounds for taking the consideration that there's moral vagueness seriously, however 'taking seriously' is construed.

Next, why then use the cumbersome notion of 'philosophical datum' and not simply say 'serious consideration'? One reason is pragmatic: by doing so, I can approach the question of how to handle moral vagueness via the framework of B, C & S-L in Chapter 5. As their framework is clear and cutting-edge, this approach is well-motivated.

Matters are more complex, however. There's a deep and general worry about using the notion of 'philosophical data.' Philosophical method might be radically different from scientific method. But B, C & S-L's stance on the former seems to emulate the latter. Like empirical scientists, philosophers should theorize based on data, they seem to say. However, does philosophy have data, understood as "starting points of inquiry" (B, C & S-L, p. 4)? There's no obvious answer, not least because the question itself is a philosophical question. Many thinkers, including philosophical heavyweights such as Hegel and Wittgenstein, puzzled their heads on the *problem of the beginning*.<sup>22</sup> This thesis is hardly the place to engage in *that* problem, and, truth be told, I am not sure how to make sense of the notion of philosophical starting points. Nevertheless, the idea that inquiry must start somewhere is intuitive and, in fact, seems analytic. Moreover, in praxis, philosophers do collect data via the methods B, C & S-L suggest (I will provide a list of these methods in section 4.3). Yet, I concede that a more careful formulation of the idea I am going to defend would have to be conditional: *if* philosophical inquiry has starting points, the consideration that there's moral vagueness counts as one of them.

Further, note that other methodological stances must also appeal to philosophical data as understood by B, C & S-L. Reflective equilibria and intellectual cost-benefit analyses also need "inputs" (B, C & S-L, p. 13). This commonality might alleviate worries about the problem of the beginning. However, it also brings us to another question: Why use B, C & S-L's method and not, for instance, Rawls's or Armstrong's?

---

<sup>22</sup> See Mácha (2022) for a recent discussion of Hegel and Wittgenstein on the problem of the beginning.

---

To make it clear: *this* chapter's argument for taking moral vagueness seriously will not depend on a specific view on philosophical theorizing. I'll just argue that the consideration that there's moral vagueness is an input for theorizing. However, the argument will depend on the idea that seemings can provide such inputs. I believe B, C & S-L would embrace that idea. And I see no principled reason why advocates of other views on philosophical theorizing should reject it. I will say more about seemings in the following sections.

In the *next* chapter (5), I will explicitly use the theorizing-related vocabulary of B, C & S-L introduced above. Presupposing that moral vagueness is a philosophical datum, I will ask how ethicists can handle moral vagueness. Can they accommodate and explain moral vagueness? Alternatively, can they argue that they don't have to do that? But while I'll explicitly use B, C & S-L's theorizing-related vocabulary, how to handle the data is a question that theorists of different stripes must answer. Proponents of other theoretical stances must also explain and accommodate moral vagueness. Therefore, using B, C & S-L's vocabulary again has a pragmatic rationale. It helps to show that moral vagueness is *hard to handle* for theorists of different stripes and, thus, must be taken seriously *in another sense*. It's not only that we should acknowledge that moral vagueness is a starting point for theorizing, but we should also acknowledge that moral vagueness is, to phrase it colloquially, a tough nut.

I hope I have successfully clarified that one doesn't have to endorse B, C & S-L's framework to the letter to follow much of my arguments in this thesis. I aim to convince philosophers of different methodological stripes that moral vagueness is a hard-to-handle input for theorizing. But what if alternative methodological frameworks do not yield that result? For instance, proponents of the cost-benefit method might not find moral vagueness hard to handle.<sup>23</sup> I believe we'd still have good reasons to take moral vagueness seriously. Let me explain.

Consider the following four options regarding philosophical methodology, which should cover sufficient ground:

**A** There's only one sound philosophical method, and we can know which it is.

---

<sup>23</sup> I'll discuss one reason for why proponents of the intellectual cost-benefit method might not find moral vagueness hard to handle briefly in 5.2.4, and further defend B, C & S-L's theoretical stance against the cost-benefit method there.

**B** There's only one sound philosophical method, and we cannot know which it is.

**C** There's a plurality of sound philosophical methods, and these different methods can produce, even if perfectly applied, contradictory end results.

**D** There's a plurality of sound philosophical methods; if perfectly applied, these different methods cannot produce contradictory end results.

Let's consider each of these options in turn.

**A** B, C & S-L defend their methodological framework by appealing to three "data about method" (Ch. 4). These data are:

- The Friendliness Datum: "Implementing a philosophical method involves engaging in such activities as: advancing arguments, raising objections, offering replies to these objections, seeking clarification, seeking explanations, displaying sensitivity to the deliverances of logic, mathematics, science, and (to a greater or lesser degree) common sense" (pp. 95f.).
- The Determinativeness Datum: "A sound method selects among multiple theories that fit the data, ruling a great many of them out of contention" (p. 96).
- The Virtues Datum: "A sound method is comprehensive, support-requiring, synoptic, multidimensional, hierarchical, and comparative" (p. 97).

Let's see their defense at work. Here's B, C & S-L's definition of the method of reflective equilibrium:

**Method of Reflective Equilibrium:** When constructing a theory about a given domain, theorists ought to attend to their considered judgments that address the central questions about the domain, seeking to achieve coherence between their considered judgments (at any level of generality) and principles that account for them through a reflective process of modification, addition, and abandonment of either the judgments or principles in case of conflict (with each other, or with any of their other relevant convictions). The best theory is the one that achieves such coherence to the highest degree relative to rivals. (p. 94)

---

Among other things, B, C & S-L argue that proponents of the method of reflective equilibrium cannot handle the Virtues Datum. More specifically, they cannot handle the consideration that a sound method must be support-requiring:

[The method of reflective equilibrium] permits theorists to introduce new claims of various sorts, so long as doing so increases equilibrium. And in fact theorists will be forced to make such additions if they are to avoid simply reprising the method's starting points (viz., the initial considered judgments). The method thus opens the door to systematic theories whose claims are unsupported by any consideration, beyond coherence, that speaks in their favor. This raises the worry that the Method of Reflective Equilibrium is not adequately support-requiring. (p. 99)

This was just one example of the strategy B, C & S-L use to defend their method as the only sound philosophical method (In 5.2.4, I will briefly present their argument against the method of intellectual cost-benefit analysis).

There is something fishy about B, C & S-L's defense, however. It is easy to see how defenders of the method of reflective equilibrium could counter. They might say, "You're theorizing about theory by employing your theory! We, the defenders of the reflective equilibrium, reject the idea that we must handle the data you impose! We would theorize about theory differently and must not follow you." How good of a response would that be? B, C & S-L could respond similarly: "You, defenders of the reflective equilibrium, theorize about theory by employing your theory! ...". Thus, we would have reached a stalemate.

It seems one would, therefore, have to find another justification for the view that B, C & S-L's theoretical framework is the only correct philosophical theory. And there's a quick argument that eludes the circularity worries just mentioned. B, C & S-L's theoretical framework is modeled on our best theory, which is scientific theory. As the 'Collect and handle the data!' imperative seems to work well in science, we have good reason to believe it will work for philosophy. But I'm not certain how many would be convinced by this appeal to B, C & S-L's emulation of scientific theory. As noted above, philosophy and science might need radically different methods. Yet, clearly, B, C & S-L's theoretical framework has much going for it. Possibly enough to have a justified belief that it is the one sound method.

**B** Maybe we should remain fundamentally skeptical about which philosophical methods are sound. In that case, B, C & S-L's theoretical framework should still be seen as a serious contender for being that theory. As just said, it can handle the data about method and is modeled on scientific theory. Moreover, it's simply a cutting-edge systematic attempt by acclaimed experts in the field, which should also lend some credence. Keeping these reasons to support B, C & S-L's theoretical framework in mind, suppose for a moment what I want to show in the following: Namely, that according to the framework of B, C & S-L moral vagueness is hard to handle. Then, a highly probable theoretical framework says: Moral vagueness is a tough nut! The consequence seems to be that we just should take moral vagueness seriously, even if we cannot know whether the framework of B, C & S-L is ultimately correct.

**C** Let's consider another possibility. Suppose there are multiple sound philosophical theoretical frameworks. Moreover, following different methods can lead to different answers to the same question. Maybe cost-benefit analysis yields the verdict that moral vagueness is easy to handle. And, if I'm right, following B, C & S-L's framework tells us that moral vagueness is hard to handle. We'd face contradicting verdicts produced by different sound methods. But if we accept this situation, we'd have to accept a methodological schism. Different schools follow different methods, which yield different results, and that's that. This would be an extreme and, for many, unsatisfactory position. Yet, for the sake of argument, suppose such a pluralism is correct. Then, nothing seems to speak against endorsing B, C & S-L's framework and its results. At least there isn't a principled way of denying that moral vagueness is a tough nut.

**D** Again, suppose there are multiple sound theoretical frameworks. But now, no matter which method you choose in pursuing a question, your path will lead to the same answer. If we are in such a situation, and different methodologies yield different verdicts, then only some of them are sound. Again, suppose moral vagueness is easy to handle according to the cost-benefit analysis and hard to handle according to B, C & S-L's framework. Some other methodologies agree with the cost-benefit analysis, others with B, C & S-L's framework. We are then back in an **A** or **B** situation, only with a pluralist flavor: we must find reasons to believe in the soundness of certain theories. And B, C & S-L's framework certainly seems a strong contender to be one of the sound theories.

In sum: There are good reasons to believe that B, C & S-L's theoretical framework is sound or among the sound ones. Even if we cannot know whether that's ultimately the case, we should hold that moral vagueness is a tough nut if B, C & S-L's theoretical framework says so. And, if B, C & S-L's theoretical framework is sound, and other sound methods yield different results, we can still endorse B, C & S-L's and maintain that moral vagueness is hard to handle.

Let me end with a final argument for following B, C & S-L's framework. One might say that B, C & S-L's framework is not an alternative to other methods but rather an overarching framework incorporating traditional methods. The method of reflective equilibrium, for instance, has an important role in what B, C & S-L call "Grounding the theory" (see their Ch. 5.2). And they praise the cost-benefit method for yielding "beliefs with a high probability given the evidence" (p. 95). The framework provided by B, C & S-L can thus be seen as a suggestion for organizing a plurality of methods. For that reason, the theoretical stance of B, C & S-L steals the 'competing' theorist's thunder. One can still do reflective equilibria and cost-benefit analyses in B, C & S-L's framework. So, many reasons in favor of those methodologies will not speak against theirs.

## 4.2 *Seemings*

This section traces the notion 'seeming' by offering a brief overview of this notion's role in recent epistemological literature.<sup>24</sup> The relevance of seemings for philosophical data becomes clear when one asks, 'How do we gather data?'. B, C & S-L (p. 63) provide the following list of procedures for data collection:

- Perception of items in one's environment
- Introspection of one's own phenomenal states
- Intuitions about thought experiments
- Linguistic judgments about the grammaticality, meaning, or truth-conditions of sample sentences
- Common sense verdicts about what exists, is known, etc.
- Imagining future possibilities
- Inductive inference from observed samples
- Sociological observations about ordinary conversations

---

<sup>24</sup> Parts of this section are based on an article I wrote for *1000 Word Philosophy* (2024).

- 
- Attending to the testimony of members of marginalized or oppressed groups
  - Statistical analyses of survey results
  - Selecting, implementing, and analyzing a model of a scenario
  - Searching for the origins of a practice through archival study
  - Interpreting a text using particular hermeneutics

As I argue in the next section, we can gather the datum that there's moral vagueness via consulting our *linguistic seemings*. While this is not a procedure that B, C & S-L mention directly, I believe they would allow for it. After all, they allow for intuitions and linguistic judgments, so it's not far-fetched to suppose they'd allow for consulting linguistic seemings. And seemings plausibly produce data in many philosophical domains. But I'm jumping ahead. This section first aims to clarify what is meant by the type of mental state called 'seeming.'

Consider the following examples: It appears to me that there's a cup of coffee in front of me and that I've had two coffees already. It's obvious that  $5 + 7 = 12$ , common sense that one shouldn't torture puppies, and the string of words 'The ran this' feels ungrammatical. We sometimes have the intuition that something's wrong with an argument or are under the impression that the time is right to take up an opportunity.

These examples illustrate that we often have a sense or feeling that something is true. Many philosophers believe that there's a type of mental attitude that is connected to such feelings, called a 'seeming.' Seemings are sometimes also described in terms of *appearances* and *intuitions*. Some philosophers, however, use 'intuition' only for intellectual seemings and 'appearance' only for perceptual seemings. So, I'll stick to the notion 'seeming' throughout this thesis to refer to the broader type of mental state. To avoid further confusion, let me also note that we do not always aim to report a mental state when we say, 'It seems to me ...'. For example, when I politely want to tell you that I disagree with your opinion, I might say, 'It seems to me you're wrong.' In the following, I'll always have reports of mental states in mind when using 'It seems ...'.

Seemings are of philosophical interest because they might explain how beliefs can be justified or reasonable without being based on other beliefs. Perhaps beliefs can be justified by how things seem to us. And seemings might explain how we could be entitled to many of our everyday beliefs. For instance, while I cannot exclude the possibility that the cup of coffee in front of me is an illusion, I don't have a good reason to take this possibility seriously. Some philosophers argue that my seeming that there's a cup of

---

coffee in front of me is enough to justify my belief that there is, indeed, a cup of coffee in front of me.

Many philosophers think that there are various kinds of seemings: the seeming that there's a cup of coffee in front of me is *visual*, the seeming that I already had two coffees is *memorial*, the seeming that  $5 + 7 = 12$  is *intellectual*, the seeming that one shouldn't torture puppies is *ethical*, and the seeming that 'The ran this' is ungrammatical is *linguistic*. However, it's up to debate whether and to what degree all those types of seemings can justify beliefs. Memorial seemings might be, for instance, generally more often mistaken than linguistic seemings.

Most philosophers think that visual appearances can provide evidence for beliefs or justify beliefs. Many philosophers believe other sorts of seemings, such as common sense and intuition, can also justify beliefs. Suppose it clearly seems to me that torturing puppies is wrong. Since I have no reason to doubt that torturing puppies is wrong, these philosophers would say that I'm justified in believing that it's, in fact, wrong to torture puppies. The view that I am at least to some degree justified to believe something when it seems true to me, and I don't have good reasons to doubt that it's true, is known as *phenomenal conservatism* (see, for instance, Huemer, 2007).

Of course, just because something seems true doesn't always mean one should believe it. Suppose there's a pencil in a glass of water in front of me. The pencil appears bent to me, but I know this to be an illusion. I know water can make straight objects appear bent. Certainly, I shouldn't believe that the pencil is bent. Some philosophers would say I'm still slightly justified in believing the pencil is bent but more justified in believing it's not. Others would say that in such a case, my seeming doesn't provide any justification. Steup (2023) and others call the view that seemings always provide at least some justification *dogmatism*. If my seeming doesn't provide any justification because I know that I'm facing an illusion, philosophers would say that my seeming is *defeated* (i.e., undercut or overridden) by my knowledge that I face an illusion.

Seemings are propositional attitudes, and as we have seen, seemings are related to beliefs. One might then wonder whether they simply *are* beliefs, perhaps just fundamental, spontaneous beliefs. Lycan (1988) thinks so. Indeed, I believe there's a cup of coffee in front of me right now, and I didn't infer this belief from another belief I had. Note that if seemings were a kind of belief, they could *not* explain how some of our beliefs can be justified without being based on other beliefs.

But consider the pencil in the water glass again. When I look at the pencil, it *merely* seems to me that it's bent, but I don't believe it is. Many philosophers hold that, therefore, we shouldn't say that seemings are just a kind of belief. One response here would be to claim that I both believe that the stick is bent and believe that it's not. This reply, however, would have the implausible consequence that I'm (deeply) irrational when experiencing such a simple visual illusion.

Another suggestion is that seemings are *inclinations* to believe, a view advocated by Sosa (2007). Inclinations to believe can result in beliefs. But we can also reject an inclination. When I see the pencil in the water, I feel drawn towards the belief that it's bent. But I don't have to accept this inclination. However, one can be inclined to believe something without having a seeming that it's true. For instance, you might lean towards believing that your favorite candidate won the elections because you wish this to be true. At the same time, it can seem to you that she lost. So seemings aren't clearly inclinations to believe either. A similar response to the one just given to the objection against the 'seemings-as-beliefs' view is given by Earlenbaugh and Molyneux (2009). You might be inclined to believe that your candidate will win and have a competing inclination to believe that she will lose.

Finally, Tucker (2013) mentions another option. Instead of a belief or an inclination to believe, a seeming might be a belief or a tendency to believe that a mental state is evidence for something.

I note that the jury is still out on whether seemings can be identified with another propositional attitude. However, for the rest of the paper, I agree with what seems to be the majority view in the relevant literature: Namely, that seemings are a particular, special type of mental attitude. Moreover, I presuppose in the following that seemings have at least *prima facie* epistemic power—at least non-overridden seemings that *p* provide some justification for believing that *p*. Finally, I presuppose a somewhat more controversial idea. If it seems to a sufficiently large number of theorists that *p* and these seemings are not overridden, theorists, considered collectively, have good reasons to believe *p*. This idea is somewhat controversial because it is usually merely assumed that seemings can provide justification for the experiencer. In any case, the claim about collective epistemic reasons via seemings is intuitive enough to take for granted in the following.

To conclude this section, let me briefly address that some philosophers doubt that seemings are a unique type of attitude. Why? Well, perhaps the alleged examples of seemings only share the name 'seeming.' DePaul (2009) expresses the suspicion that

there's no unique type of mental state to which the name 'seeming' refers. As already noted, the examples given above are pretty diverse. Pointing out that every mental state we call 'seeming' has the common feature of typically producing inclinations to believe (while seemings are not identical to inclinations to believe) would not convince skeptics of the existence of a unique mental state called 'seeming.' After all, wishful thinking and other mental states also typically lead to inclinations to believe. However, I believe that the admittedly somewhat obscure notion of 'feeling of truth' is enough to appeal to as an answer to this worry. Every mental state we call 'seeming' has the common feature of coming with a feeling of truth.

Another reason to doubt that there are seemings is this: you're probably aware of many mental states, including beliefs and inclinations to believe. But can you find seemings in your mind? If not, you aren't alone. Famous philosophers haven't found them either. Conee (1988) and Williamson (2007) report that they are unaware of having seemings. On the other hand, throughout history, many philosophers have written about attitudes that are at least very similar to seemings, using names such as *phantasia* or *clear and distinct perception*.<sup>25</sup> This longstanding interest in such mental states gives some confidence that seemings exist even though they're hard to grasp.

In any case, it certainly seems to me that there are seemings. Let me now say a few more words about *linguistic seemings* in particular.

### 4.3 *The Seeming That There's Moral Vagueness as a Linguistic Seeming*

Linguistic seemings are linguistic because they are directed at a proposition that has an application of a property to one or more linguistic objects as its content. Here are some such propositions:

- The string of words 'The ran this' is ungrammatical.
- The predicate '... is a flying horse' has an empty extension.
- 'Serendipity' sounds good.
- '... is a bachelor' is synonymous with '... is an unmarried man'.

---

<sup>25</sup> Aristotle uses the word *phantasia* (for instance, in *De Anima*, [350 BC] 1986), and Descartes uses the phrase *clear and distinct perception* (for instance, in *Meditations on First Philosophy*, [1641] 1993). McAllister (2023) provides more historical examples.

- The sentence ‘Çok teşekkür ederim’ is understandable to people who can speak Turkish.
- The phrase ‘Could you please send the email as soon as you can? Thanks!’ is preferable to ‘Send the mail asap!’ in this context.
- ‘lpink’ is not a possible combination of sounds in English.
- ‘bank’ is ambiguous.

The listed propositions are propositions one can have seemings about.<sup>26</sup> To me, at least, it does seem that the string of words ‘The ran this’ is ungrammatical, that the predicate ‘... is a flying horse’ has an empty extension, and so on.

Yet, some might reply that what I call a linguistic seeming is not a real seeming. In contrast to perceptual seemings, the reply might go, what I call linguistic seemings are just plain beliefs. Some might add that linguistic seemings are not even fundamental beliefs. What I call linguistic seemings are beliefs that must be based on other beliefs one has (for instance, beliefs about grammar, horses, bachelors, and so on).

The linguistic phenomenon of ‘Escher-Sentences’ (see, for instance, Townsend & Bever, 2001) presents cases that hopefully convince those initially skeptical about linguistic seemings. Have a quick look at the following proposition:

The sentence ‘More people have been to Russia than I have’ is grammatical.

I confess that ‘More people have been to Russia than I have’ appears grammatical to me even after I have judged it not to be. So, we have a linguistic case analogous to the bent-pencil illusion. Therefore, we have a case where it’s plausible that we do not have two contradictory beliefs but a seeming and a conflicting belief. It’s not the case that I believe that the sentence ‘More people have been to Russia than I have’ is ungrammatical *and* that it’s grammatical. Yes, I believe it’s ungrammatical, but it seems to me that it’s grammatical. I think this would be a common reaction when asked to reflect on ‘More people have been to Russia than I have.’ Alternatively, one might claim that I *do* have a

<sup>26</sup> As the list shows, linguistic seemings can be further subdivided into syntactical, phonological, semantical, sociolinguistic, pragmatical seemings, and so on. It is not necessary to do that here, however.

conflicting, somehow special belief in such illusion cases. As long as such a special belief can be called ‘seeming’ and has justificatory power, I’m OK with that.

Moreover, appealing to linguistic seemings, or linguistic intuitions as they are more often called, is a standard praxis in linguistics (see Maynes and Gross, 2013). Of course, linguists gather data in many other ways. They advert to corpora of actual usage, they perform reaction-time experiments and fMRI, they compare across languages, and so on. But appealing to raw intuitions is an indispensable praxis in linguistics.

That said, let us assume that there are linguistic seemings and move on. Here’s the crucial idea of this section. An important class of linguistic seemings consists of seemings that some linguistic item is in the extension of a predicate. Let me explain.

First, note that the sentential pendants to the propositions listed above are predications of predicates to linguistic items. ‘... is ungrammatical’ is predicated on the string of words ‘The ran this.’ ‘... has an empty extension’ is predicated of the predicate ‘... is a flying horse’. (Thus, ‘... has an empty extension’ can be called a *second-order predicate*). ‘... sounds good’ is predicated on the word ‘Serendipity.’ And so on.

Next, note that one need not fully grasp the concept of *Ungrammaticality* to claim that any theory of grammar has to take into account that the string of words ‘The ran this’ is ungrammatical. Nor does one have to fully grasp *Horseness* and *Extension* to claim that any relevant theory must consider that ‘... is a flying horse’ has an empty extension. And one doesn’t have to fully understand *Good-Soundingness* to claim that aestheticists must consider that ‘Serendipity’ sounds good.<sup>27</sup>

What this shows is that one can have linguistic seemings about linguistic items without fully understanding what the corresponding concepts consist in. But let me go over this in more detail by returning to the question of whether moral predicates are vague.

‘... is vague’ is a second-order predicate in the sense that it must be applied to terms that refer to predicates. It is in this respect like ‘... has an empty extension’ (or ‘... is synonymous to the predicate ...’ or ‘... is a one-place predicate’, and so on).<sup>28</sup> I believe

<sup>27</sup> If one were to claim the opposite, one would, as Geach (1966) famously called it, commit the ‘Socratic Fallacy’.

<sup>28</sup> “... is a moral predicate” is also second-order. So, the verdict that moral predicates are vague is actually a verdict that involves two second-order predicates. I don’t think this has any important implications, so I leave the matter there.

---

we can have a linguistic seeming that moral predicates are in the extension of the predicate ‘... is vague’ without understanding the concept *Vagueness*.

To avoid confusion, let me explain how I use some notions. Let me use the notion *concept* in a way that is neutral towards specific semantics: a concept is just what a predicate *expresses*.<sup>29</sup> Let’s call what a concept *refers to a property*, staying neutral on the metaphysics of properties.<sup>30</sup> I use *extension* to refer to the set of subjects a predicate can be truthfully applied to. Finally, I say that quotations *mention* the expression that they quote. For instance, the quotation ‘Kaj’ mentions the term ‘Kaj,’ and the quotation ‘... is a great upcoming philosopher’ mentions the predicate ‘... is a great upcoming philosopher’.<sup>31</sup> With that being said, let me sketch the semantics of ‘... is morally permissible’ is vague.’

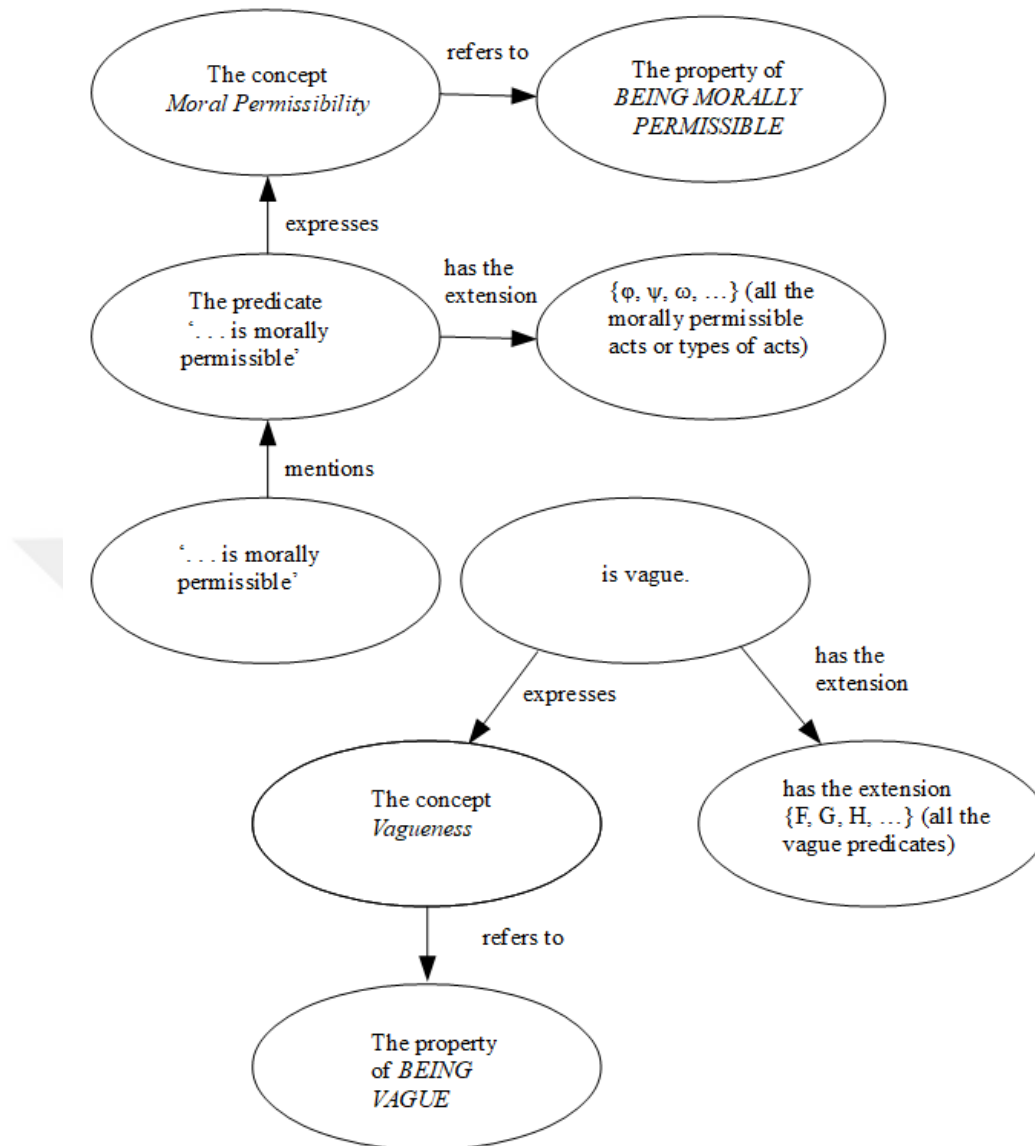
---

<sup>29</sup> I use ‘expresses’ and ‘concept’ in a broad, theory-neutral sense. Fregeans might think of expressing as the relation between an expression and its correlating abstract object; use-theory advocates might think of expressing as the relation between an expression and its role in human social behavior; truth-condition-theory advocates might think of expressing as the relation between an expression and its role in a sentence’s truth function; and so on. In general, the term ‘expresses’ is just meant to be a term for the relation between an expression and what it expresses. Whatever an expression expresses, I call ‘concept’.

<sup>30</sup> Properties might, for instance, be *dispositional* and thus have necessary connections with other entities or *categorical* and thus have no necessary connections with other entities (see Bird, 2007). Whether moral and linguistic properties are dispositional or categorical appears to me to be especially puzzling. However, these conundrums I flag and leave aside.

<sup>31</sup> One might also say that ‘Kaj’ *refers* to the term ‘Kaj’ and that ‘... is a great upcoming philosopher’ *refers* to the predicate ‘... is a great upcoming philosopher’. To avoid confusion about the term ‘refer’, I here use the common notion of mentioning.

Figure 1: The Semantics of Moral Vagueness



The sentence ‘... is morally permissible’ is vague’ contains a quotation (‘... is morally permissible’) and a predicate (... is vague). The quotation ‘... is morally permissible’ mentions the predicate ‘... is morally permissible’. This predicate has the set of all morally permissible acts or types of acts as its extension and expresses the concept *Moral Permissibility*. This concept refers to the moral property of *BEING MORALLY PERMISSIBLE*. The predicate ‘... is vague’ has the set of all vague predicates as its extension and refers to the concept *Vagueness*, which refers to the linguistic property of *BEING VAGUE*.

Here’s why I’m detailing the semantics of ‘... is vague’. Ask yourself how one can *know* whether something is in the extension of a predicate. It’s natural to assume that one

has to understand the concept the predicate expresses to do so. In our context, this would mean that we would first have to understand *Vagueness* to know whether moral predicates are vague. But, as the characterization problem shows, we don't understand the concept *Vagueness*. What's worse, the concept of *Vagueness* has to do with the semantics of first-order predicates. To know whether a predicate is vague, one might have to understand *how* a first-order predicate expresses a concept (or multiple concepts). Additionally, one might have to understand *how* it refers to a property (or numerous properties). So, we would also have to understand *Moral Permissibility* to know whether '... is morally permissible' is vague.

However, one doesn't always have to fully grasp a concept to know that a consideration involving this concept is relevant. To reiterate, one doesn't have to go down the route via concepts to have the seeming that a linguistic item is in the extension of a predicate. It can seem that 'The ran this' is ungrammatical absent a full understanding of *Ungrammaticality*. One does not have to fully grasp *Horseness* and *Extension* to have the seeming that '... is a flying horse' is in the extension of '... has an empty extension'. It seems to me that 'Serendipity' is in the extension of '...sounds good'. But I couldn't tell you what *Good-Soundingness* consists in. What I—and I believe most of us—find it natural to think is that these seemings are enough to warrant serious considerations. In other words, they produce data for the corresponding theories. In fact, arguing that one must first understand the concept to take the consideration produced by these seemings seriously would have things backward. One doesn't have to understand *Ungrammaticality* to take 'The ran this' is ungrammatical' seriously. The reverse sounds right. The proposition <'The ran this' is ungrammatical> is the input for a theory of grammar, aiming for an understanding of concepts like *Ungrammaticality*.

Now consider the following propositions:

- '... is completely behind the table' is not vague.
- '... is far away' is vague.
- '... is exactly 180cm tall' is not vague.
- '... is tall' is vague.
- '... is a prime number' is not vague.
- '... is a high number' is vague.
- '... is a child' (in the sense of being a first-generation descendent) is not vague.
- '... is a child' (in the sense of being of a certain age) is vague.

- ‘... is a grandmother’ is not vague.
- ‘... is old’ is vague.
- ‘... is a bachelor’ is not vague.
- ‘... is introverted’ is vague.
- ‘... is fifty meters long’ is not vague.
- ‘... is long’ is vague.
- ‘... arrived at exactly three o’clock’ is not vague.
- ‘... arrived early’ is vague.
- ‘... were precisely one hundred people’ is not vague.
- ‘... were around hundred people’ is vague.
- ‘... is exactly in the center’ is not vague.
- ‘... is in the center’ is vague.

The listed propositions have the (non-)application of the linguistic property of *BEING VAGUE* as their content. The point of this list is to show that we have linguistic seemings that concern vagueness. These seemings are independent of whether we can say what vagueness consists in. I, for instance, have no clear conception of what vagueness is, yet all listed propositions seem true to me. When considering these propositions, I have a feeling of truth. This feeling makes me inclined to believe them.

After all these preliminaries, now think again about Cheryl, who is considering an on-demand abortion, and Darryl, who is watching his two-year-old daughter play in a city park. Doesn’t the following verdict seem true?

‘... is morally permissible’ is vague.

It will seem true to many people that ‘... is morally permissible’ is vague, independent of theoretical convictions about vagueness and moral permissibility. One doesn’t have to fully understand *Vagueness* and *Moral Permissibility* to feel that ‘... is morally permissible’ is vague. This will be the case for many, even when their theoretical convictions tell them there’s no moral vagueness. You might, for instance, believe that vagueness entails soriticality and that ‘... is morally permissible’ is not soritical. Yet, you might still feel that ‘... is morally permissible’ is in the extension of ‘... is vague’. Or, you might strongly hold on to the view that moral predicates qua moral predicates cannot

be vague yet nevertheless feel the pull of the consideration that moral predicates are vague.

To many philosophers, it indeed seems that there's moral vagueness. Constantinescu (2014) argues that moral predicates have all the alleged features of vagueness and are thus vague. Yet, revealingly, he says:

So far, I have simply assumed without argument that moral predicates are vague . . . *I find this claim absolutely intuitive* . . . I also hope to convince those who . . . find this idea doubtful . . . I shall therefore give five interrelated reasons for thinking that at least some (and probably most) of our moral predicates are vague. (2014, pp. 156f., my emphasis)

In addition, he claims that in contrast to the Pythagoreans, who allegedly assigned the number 4 to justice,

Aristotle *sounds much more sensible* when, reflecting upon the subject matter of his inquiry in the Nicomachean Ethics, he advises that “we must be content, in speaking of [fine and just actions, and of goods in general], to indicate the truth roughly, . . . for it is the mark of an educated man to look for precision in each class of things just so far as the nature of the subject admits” (1094b, 19–25). *Common sense* appears to side with Aristotle here: the very attempt to assign cardinal numbers to essentially imprecise, indeterminate values like justice, happiness, goodness, and the like, strikes most of us as deeply incongruous. (2014, p. 152, my emphasis)

The gist of what Constantinescu wants to say is that it is sensible or common sense that predicates like ‘... is just’, ‘... is happy’, or ‘... is good’ are vague. In other words, it seems to him that there's moral vagueness.

Hawthorne (2022, p. 220) reports that cases like Darryl's seem to be cases of moral vagueness:

Let us take what seems like a paradigm case of moral vagueness . . .:

Diversions: Darryl is watching his two year old daughter play in a city park. It is permissible to divert his attention from her for 1 second. It is not permissible to divert

his attention from her for 5 minutes. Is it permissible to divert his attention for 30 seconds? 31? 32? . . .

This certainly seems like a case of vagueness.<sup>32</sup>

Next, here's how Dougherty (2014, p. 353) introduces moral vagueness:

It *seems* that you may save your friend at the expense of saving a few strangers, but not at the expense of saving very many strangers. How many is very many? No precise number *appears* to fix a threshold. The boundary *seems* fuzzy, and correspondingly some actions, for instance saving your friend rather than thirteen strangers, are borderline cases of a morally permissible action. Thus the term “morally permissible” is vague. Similarly, there *seems* no precise number of milliseconds of pleasure that would compensate someone for a broken pinky toe; nor does there *seem* to be an exact maximal number of cents that you must spend on a taxi in order to keep your promise to meet someone at a restaurant at 7 pm. (My emphasis)

Dougherty here appeals to straightforward arguments from boundarylessness and borderlineness. Note, however, how often Dougherty mentions seemings and appearances. I think this indicates that it seems to Dougherty that moral predicates are vague. Yes, he reports seemings about alleged features of vagueness rather than the raw linguistic seeming that moral predicates are vague. But he might just do that because he wants to appear more argumentative. In light of the characterization problem, appealing to raw linguistic seemings would have been the better strategy.

Sud (2019, pp. 684f.) also appeals to seemings of plausibility and appearances:

A sorites series of hair loss makes it *plausible* that some people are borderline bald. Similarly, sorites series like those described in *Amputations*<sup>33</sup> make it *plausible* that some actions are borderline wrong . . . initial reflection on the sorites series of hair

<sup>32</sup> Hawthorne then goes on to argue that we might not want to trust this seeming—that is, he attempts to *disable* the datum that there's vagueness. I will address a disabling-strategy similar to Hawthorne's in 5.2.

<sup>33</sup> Sud is here referring to a case originally given by Schoenfield (2016, p. 263; see my Introduction).

---

loss and the sorites series of amputations makes baldness and wrongness *appear* analogous in terms of their vagueness. (My emphasis)

Sud's reasoning is not a straightforward argument. His first claim is that both predicates that are usually seen as being vague (such as '... is bald') and moral predicates are soritical. His second claim is this: because both predicates that are usually accepted to be vague and moral predicates are soritical, it *seems plausible* that both types of predicates have borderline cases. His third claim is that because of this borderliness, predicates that are usually accepted to be vague and moral predicates are *apparently* analogous. Again, it is interesting how often Sud appeals to *felt plausibility* and *appearances*. It is even more interesting that Sud realizes that borderliness does not sufficiently characterize vagueness. As mentioned above, Sud understands borderline cases as "cases of indeterminacy" (2019, p. 685) but notices that borderliness does not imply vagueness.

The upshot of the characterization problem is that philosophers would do better by appealing to the raw linguistic seeming that moral predicates are vague. I suppose this seeming underlies their reasoning anyway.

Let me recapitulate what I believe I have achieved in this chapter until now. I hope I have made it plausible that we can gather the datum that there's moral vagueness via linguistic seemings. Seemings produce data in many domains and are widely appealed to in linguistics. We do have seemings about whether a predicate is vague or not. And when considering cases like Cheryl's or Darryl's, it will appear to many that '... is morally permissible' is vague. Given the characterization problem, we cannot do more than appeal to raw linguistic seemings.

Let me end this chapter by reflecting on what it's like to have the seeming that a predicate is vague. Hopefully, these reflections further support the claim that we can gather the datum that there's moral vagueness via seemings. I will defend the view that moral vagueness is a datum that is quite difficult to handle in the following chapter.

#### **4.4 The Phenomenology of Vagueness**

Here, I claim that vagueness has a distinct phenomenological profile. Think of the predicate '... is bald'. Now think of the predicate '... is prime'. Certainly, there's a felt

contrast between the two predicates: they seem semantically different.<sup>34</sup> This section aims to describe this contrast.

The first phenomenological observation I'd like to point out is that when reflecting on a vague predicate, we are confronted with unclarity about its application conditions. We are somewhat in the dark about what can and cannot be slotted in the gap of '... is bald'. In contrast, non-vague predicates appear to have clear application conditions. One might be in the dark about *what* applies to '... is prime' (does, for instance,  $2^{82,589,933}-1$  apply?). Yet, it certainly feels like there must be a determinate answer. The function of '... is prime' feels clear. It maps every number to either the extension or anti-extension. Of course, one might respond that in the case of '... is prime', we simply know the function. But even if we don't know a predicate's function, the predicate can appear to have a clear function. Vagueness, however, comes with feelings of unclarity of application conditions. But this is not yet a distinct characterization of the phenomenology of vagueness. Other predicates share such unclarity of application conditions.<sup>35</sup>

As mentioned in 3.1.1, Fine (1975) illustrates vagueness through the predicate '... is nice<sub>1</sub>'.  $n$  is nice<sub>1</sub> if  $n > 15$  and  $n$  is not nice<sub>1</sub> if  $n < 13$ . So, '... is nice<sub>1</sub>' has unclear application conditions because 14 is neither mapped to its extension nor its anti-extension. Against Fine (but with Sainsbury (1996) and many others), I think a predicate like '... is nice<sub>1</sub>' is not vague. One should rather call such predicates *underdetermined* (or if you prefer, *incomplete*).

The what-it's-likeness of vagueness is distinctly different from the what-it's-likeness of underdetermination. Many authors shed light on this difference. Let's first hear Wright (2001, p. 70) talking about what I call *vagueness-related borderline cases* (in contrast to underdetermination-related borderline cases).

Borderline cases come across as hard cases: as cases where we are baffled to choose between conflicting verdicts about which polar verdict applies, rather than as cases which we recognize as enjoying a status inconsistent with both . . . the indeterminacy

---

<sup>34</sup> Of course, every predicate will seem different from every other predicate just because it expresses a different concept, has a different extension, and so on. I here mean that '... is bald' and '... is prime' seem different apart from those trivial differences.

<sup>35</sup> To further support that there's a phenomenological difference between cases of vagueness and cases of ignorance, note that in the empirical study of Alxatib & Pelletier (2011) participants almost never answer "I don't know" to the question of whether a borderline case applies to its predicate.

will be initially manifest not in (relatively confident) verdicts of indeterminacy but in (hesitant) differences of opinion (either between subjects at a given time or within a single subject's opinions at different times) about a polar verdict, which we have no idea how to settle — and which, therefore, we do not recognize as wrong.

Very similarly, Constantinescu (2014, p. 153) writes:

In a borderline case, the semantic criteria governing a predicate's application seem to pull in opposite directions: if  $a$  is a borderline instance of the predicate " $P$ ," then there are strong reasons in favour of classifying  $a$  as  $P$ , and strong reasons in favour of classifying  $a$  as not- $P$ . The clash between such reasons can be experienced internally by one speaker, leading to perplexity or ambivalence, or externally by two or more speakers, leading to disagreement.

Schiffer (2020, pp. 156ff.) describes the phenomenology of vagueness-related borderline cases like this. Imagine you are confronted with a borderline case of a bald person. Then, "our state of mind will be one of ambivalence, in that you will have some inclination to say that [the individual in question] is bald and some inclination to say that he's not bald."

Finally, Sorensen (2011, p. 237) notes "continuity, ambivalence, and indecision" as the "stereotypical features of vagueness."

Neither of these descriptions fits *underdetermination-related borderline cases* like 14 in relation to ' $\dots$  is nice<sub>1</sub>'. I'm not baffled to choose between whether ' $\dots$  is nice<sub>1</sub>' applies to 14 or not. I do not experience differences of opinion about whether 14 is nice<sub>1</sub>. I don't feel that the semantic criteria of ' $\dots$  is nice<sub>1</sub>' pull in opposite directions. I have *no* reason to classify 14 as nice<sub>1</sub> or as not-nice<sub>1</sub>. I'm not ambivalent. I am neither inclined to say that 14 is nice<sub>1</sub> nor that 14 is not nice<sub>1</sub>. ' $\dots$  is nice<sub>1</sub>' doesn't feel continuous. I'm not indecisive but rather *indifferent* about whether 14 is nice<sub>1</sub>. 14 sits, so to say, comfortably between categories. A man with 1856 hairs on his head sits awkwardly between the clearly bald and the clearly hirsute.

Now, of course, we can feel ambivalent about many things. Here's a list of examples from Razinsky (2016, p. 3).

Reuben wants to go to the theater, but also to stay at home and rest. Simeon wants and does not want to be respected. Levi loves his mother, but hates her as well. Judah loves his father, yet is a bit ashamed of him. Zebulun thinks that the show is fantastic, although, perhaps, too long. Issachar judges Judah to have been right, and yet terribly wrong, in convincing his brothers to sell Joseph. Dan understands that the family should leave for Egypt, but really does not want to. Gad wants to harm Joseph, for which he despises himself. Asher wants to save Joseph, but does not try to. Naphtali does not know what he wants. While ordering the arrest of his brothers, Joseph is sure that he loves them. Grieving for his dead mother for years and years, Benjamin is ambivalent as to whether it is good that he feels this way. Dinah, having heard that her brothers sold Joseph, cannot believe it, yet believes it anyway.

Some of Razinsky's examples might illustrate the ambivalence we feel when confronted with a vagueness-related borderline case. Zebulun's ambivalence might arise from being confronted with a vagueness-related borderline case of '... is a good show'. However, many of Razinsky's examples seem to rather exemplify ambivalence produced by normal uncertainty. Reuben, for instance, might just be unsure about whether going to the theatre or staying at home is better. He might have a 0.5 (standard) partial belief for each option. What might be called *ordinary-uncertainty-ambivalence* has, however, a different phenomenological profile than the ambivalence that vagueness produces. Ordinary-uncertainty-ambivalence does not feel unsolvable. Here's how Schiffer puts it.

To regard something as a [vagueness-related] borderline case of a property is to have some inclination to judge it to have the property, and some inclination to judge it not to have the property, when it's clear that nothing can be forthcoming to resolve the matter for one (all relevant facts are already in). (2003, p. 5)

What's special about vagueness-related ambivalence? In the remainder of this section, I explore this question. I suggest the following, probably not completely sufficient, phenomenological distinction of vagueness: We typically feel what I call *resemblance-ambivalence* about vague predicates.<sup>36</sup> A borderline case of '... is bald' is,

---

<sup>36</sup> Some visual illusions such as the infamous rabbit-duck might be cases of resemblance-ambivalence without being vagueness-related borderline cases. In some sense, the rabbit-duck seems 'duck-ish' and not

as one might say, *baldish*: he seems bald; he's similar to clear cases of baldness; he's *sort* of bald. At the same time, he's not *baldish*; he seems not-bald; he's similar to clear cases of being not bald; he's *sort* of not-bald. In contrast to ordinary-uncertainty-ambivalence, I construe resemblance-ambivalence to be irresolvable.

Here's a conceptual reason to believe we're typically resemblance-ambivalent about borderline cases. Borderline cases come in degrees. These degrees plausibly correspond to how similar the borderline case is to clear cases. If you pluck hair after hair from a clearly not-bald man, he'll be borderline bald at some point but still appear somewhat not-bald. After more plucking, he'll be less similar to not-bald men, become more similar to bald men, and so on. These competing similarities give us reason to categorize the man as bald and as not-bald. We're ambivalent about whether a borderline case of *F* is *F* or not *F* because the case is similar to *F*s and *not-F*s.

We also have good empirical reasons to believe that resemblance-ambivalence, at least to some extent, captures the phenomenology of vagueness for competent language users.

Multiple studies support the idea that order effects occur when categorizing borderline cases. It seems well-established that how competent language users classify borderline cases is influenced by the order in which they are presented with cases. It makes a difference whether you're first presented with a clearly bald person, then a slightly less bald person, and so on, until you are confronted with a borderline case of baldness, or whether you're first presented with a clearly not-bald person, then a slightly less not-bald person, and so on, until you face a borderline case of baldness. And it makes a difference whether you see a borderline case of baldness next to other cases.

The results of Egré et al. (2013) show the order effect of *enhanced contrast*, which was already observed by Kalmus (1979). Subjects were shown either a clearly orange patch, then a slightly less orange patch, and so on, until they reached a clearly yellow patch or the other way around. To illustrate the results, suppose there are thirty color patches, blending from clear orange (patch #1) to clear yellow (patch #30). If we start from patch #1, we might categorize patch #15 as yellow because we recognize a sufficient

---

'duck-ish', or, maybe more precise, *sort* of representing a duck but also *sort* of not representing a duck. The questions of whether vagueness-related resemblance-ambivalence is of the same type as visual-illusion-related resemblance-ambivalence, and of whether there are other forms or cases of resemblance-ambivalence are, I believe, very big issues, doing full justice to which goes beyond the scope of this thesis.

similarity to yellow patches; it is yellowish. But if we start from patch #30, we might already categorize patch #21 as not-yellow because we recognize a sufficient similarity to not-yellow patches; it is not-yellowish. We usually don't switch categories again after doing so. So, if we start from patch #30, patch #15, categorized as yellow before, will now be classified as not-yellow. The same color patch that appeared sufficiently yellowish to count as yellow also appears sufficiently not-yellowish to count as not-yellow. This phenomenon of enhanced contrast has been replicated by Raffman (2014) and Stöttinger et al. (2016). Generally speaking, the findings of Kalmus, Egré et al., Raffman, and Stöttinger et al. support the view that borderline cases seem *F*-ish and not-*F*-ish to us.

Raffman (2014) not only largely replicated the phenomenon of enhanced contrast but also observed another order effect: *hysteresis*. Again, suppose thirty color patches blend from clear orange (patch #1) to clear yellow (patch #30). Raffman reports that when the order of patches is reversed after competent language users shift categories, they tend to stick to their new categorization—that is, if we are first shown patch #1, then patch #2, and so on, we will at some point switch categories to yellow (or borderline) but continue to apply that category, even when we're returning to the patches we already saw. Let's again say we switched at patch #15 to yellow. If we are then shown again patch #14, we tend to also categorize it as yellow, even though directly before we did not do so. Most subjects in Raffman's experiment stuck to their new category for more than one patch after the order was reversed. Again, this experiment supports the view that borderline cases seem both *F*-ish and not-*F*-ish to us. Raffman furthermore says that because the phenomenon of hysteresis is not restricted to the perceptual domain<sup>37</sup>, we have good reasons to expect hysteresis also in non-perceptual vague predicates like '... is bald'.

In addition, Hampton et al. (2005) report a *contrast effect* for borderline cases. In their main experiment, subjects were first presented with several color patches on a spectrum. It was determined which color patch the subjects individually categorize as equally similar to both opposing colors, which plausibly means that this patch is considered a borderline case by the subject. After that, subjects were asked to categorize two color patches side by side. When the borderline color patch was seen next to increasingly clear cases, the probability of subsuming the borderline case under that category decreased. There's only a tendency to categorize both patches as the same color

---

<sup>37</sup> See, for instance, Oliva et al., (1992), Tesser & Achee (1994), Van Der Maas et al., (2003).

---

when the borderline patch is seen next to a very similar patch. Let's say that patch #15 appears borderline to you; it's orange-yellow. If you see #15 next to #16, there's a relatively high probability that you claim both are yellow. If you see #15 next to #14, there's a reasonably high probability that you categorize neither of them as yellow. But the further away you move, the less likely you'll be to put both cases in the same category. Plausibly, we here see categorizing borderline cases based on similarity (or lack thereof) in action. Hampton et al. replicated the contrast effect for color categorization with morphs presenting ambiguous pictures between a dog face and a cat face.

The empirical evidence I've just summarized suggests that the phenomenology of borderline cases involves resemblance-ambivalence, at least when we must categorize borderline cases. I conclude that when a predicate seems vague to us, we have a feeling that it has such cases.

The aim of this section was to offer some reflections on *how* vague predicates seem to us. Now, consider again the cases of Cheryl and Darryl. Doesn't it seem that abortion after 151 days is *permissibleish* and *not-permissibleish*? Doesn't it seem that looking at your phone for 30 seconds is awkwardly between categories, as well similar to permissible acts and to impermissible acts?

Of course, after this long and winding road, you might still simply say, "I'm not having such seemings." You might say that moral predicates do not make you feel resemblance-ambivalence. It might seem to you that there could be something forthcoming to resolve the matter of whether Cheryl is permitted to abort—you might just feel ordinary-uncertainty-ambivalence. Or you might not feel any ambivalence. The what-it's-likeness of moral predicates might be to you rather like the what-it's-likeness of underdetermined predicates. That's all right for me. My goal was to offer an alternative route to arguing that theorists, *considered collectively*, should seriously consider moral vagueness. The foregoing reflections contribute to clarifying how vagueness seems to us. But it's not up to you or me whether the route via seemings establishes that moral vagueness is a philosophical datum. That depends on whether enough theorists share the seeming that there's moral vagueness. What I attempted to achieve in this chapter is merely to pave the way to understanding what this seeming is.

## Chapter 5: HANDLING MORAL VAGUENESS

In the preceding chapters, I discussed whether moral vagueness is a consideration that theorists must take seriously. Seemings can produce such considerations. I offered clarifications about what it means to have the seeming that there's moral vagueness. Maybe it's already the case that it seems to the majority of ethicists that there's moral vagueness. If not, the clarifications I offered might help bring this about.

This chapter shows what it would mean to take the consideration that moral predicates are vague seriously—that is, what the discursive consequences are when moral vagueness is a datum. Recall: Theorists must handle the data, and they can do so in a positive or negative way. The positive way is to accommodate or explain the data. In this final chapter, I focus on the negative ways moral theorists could handle the datum of moral vagueness. The reason for this is, first of all, that I'm interested in highlighting why moral vagueness is important. I will do so in the following by showing how difficult it is to handle moral vagueness in a negative way. This shows that if moral vagueness is a datum, it is a datum that is not easy to manage. Secondly, I am interested in discussing moral vagueness as a datum for moral theories. And I'm simply not sure how moral theories could handle moral vagueness in a positive way. One might say that moral theorists could accommodate moral vagueness by showing that their theory coheres with strong theories about vagueness. Imagine metaethical theory *A* runs more smoothly with a solid theory of vagueness than metaethical theory *B*. Proponents of *A* then seem to be able to claim that moral vagueness is more likely to exist given *A* than *B*. However, this strategy is foremost a negative way of handling moral vagueness: Namely, by passing on the job of explaining moral vagueness to theorists of vagueness.<sup>38</sup>

Handling moral vagueness via passing on will be the topic of section 5.1. Section 5.2 discusses handling moral vagueness via offering data-disablers of the first kind. There, I discuss the armada of tests that Kultgen suggests in his “Vagueness-Last” approach. In my view, these tests do not suffice to show that the seeming that there's moral vagueness is faulty. In section 5.3, I discuss a data-disabler of the second kind—that is, a consideration that seems to strongly contradict the datum that there's moral vagueness.

---

<sup>38</sup> A worry here is that general theories of vagueness might fail to explain why moral predicates in *specific* are vague and that philosophers of vagueness would just re-defer this datum to the ethicists. This worry, in fact, points to a larger methodological issue not discussed by B, C & S-L: how to handle cross-domain data? Unfortunately, I have too little space to pursue this fundamental matter further.

The consideration in question is Chang's claim that we can stipulate vagueness-related borderline cases away.

What about the fourth way of negatively handling data? Let me confess that I can't see how pleading unexplainability in front of moral vagueness could be anything but a last resort. It might be that some theorists will seek refuge in the 'It's a brute fact' defense. But I shall look past this way of handling moral vagueness in the following discussion.

### **5.1 *Passing On Moral Vagueness***

Theorists must handle the data. In our case, ethicists must handle moral vagueness. The topic of this section is whether stance-independent moral realism (henceforth just moral realism) can do so by passing on the datum of moral vagueness to philosophers of vagueness.

In the following, I shall propose that moral realists cannot *easily* pass on the datum of moral vagueness. Why focus on moral realism? It has been argued that moral vagueness somehow threatens other metaethical theories as well. Baima (2014) argues that metaethical non-cognitivists are in trouble if moral predicates are vague. Dougherty (2014) raises problems for stance-dependent realists. Kultgen points out that any theory that identifies moral terms with precise natural terms is at odds with moral vagueness (2018, p. 53). Moreover, some first-order ethical theories arguably rule out moral vagueness from the outset. Schoenfield briefly mentions scalar theories as candidates for 'vagueness-free' theories since such approaches lack moral on/off terms (2016, p. 264). Yet, as Dunaway (2017, p. 41) argues, predicates like '... is 0.6 permissible' might also display symptoms of vagueness.

However, the main discussion about the metaethical implications of moral vagueness focuses on how moral realism can handle moral vagueness. Moreover, this discussion can be helpfully examined through the lens of B, C & S-L's framework. This is so because it centers on whether moral realism is compatible with certain theories of vagueness. And showing compatibility with theories from another domain is clearly a way, or maybe a prerequisite, of passing on a datum to such other theories. To pass on a datum to a theory from another domain, one has to show that one's theory coheres with that other-domain theory. In this section, I identify several problems that moral realists face when attempting to show that moral realism coheres with theories of vagueness.

It is commonly assumed that theories of vagueness fall into one of three categories:

*Onticism* Vagueness would remain even if we spoke a perfect language (a language that would contain all and only predicates needed to describe reality completely) and if we were omniscient.<sup>39</sup>

*Epistemicism* Vagueness would vanish if we were omniscient speakers of our actual language.<sup>40</sup>

*Semanticism* Vagueness would only vanish if we were omniscient speakers of a perfect language.

According to onticism, vagueness is out there in the world. A proposition like <Someone with 1856 hairs on their head is bald> is indeterminate because there is no corresponding determinate state of affairs. Usually, onticists claim that predicates are vague because the properties they refer to are vague and that those properties exist independent of our stances. Shafer-Landau (1995) and Schoenfield (2016) argue that moral realists should adopt that view. Their main argument for this view is that neither semanticism nor epistemicism is unproblematic for moral realists.

Yet, onticism is itself problematic for moral realists. To see why, let me adapt an argument from Barnett (2009), assuming that it's indeterminate whether looking at your phone for 30 seconds is permissible.

**P1** It is not determinate that looking at your phone for 30 seconds is permissible, and it is not determinate that looking at your phone for 30 seconds is impermissible.

**P2** If it is not determinate that  $p$ , 'p' is not true.

**From P1, P2:** 'Looking at your phone for 30 seconds is permissible' is not true, and 'Looking at your phone for 30 seconds is impermissible' is not true.

**P3** If 'p' is not true,  $\sim p$

**From P1, P2, P3:** It's not the case that looking at your phone for 30 seconds is permissible, and it's not the case that looking at your phone for 30 seconds is impermissible.

<sup>39</sup> Sud gives a more subtle definition (2019, p. 690).

<sup>40</sup> Sorensen's epistemicism (1988, 2008) doesn't fit this characterization and is underexplored in the debate.

It seems to follow that it's both true and false that looking at your phone for 30 seconds is permissible, and that conclusion will be noxious to moral realists, who typically hold that moral propositions cannot be simultaneously true and false.<sup>41</sup>

Of course, there's room for skirmishing here. For instance, one might reject P2 by denying that truth is disquotational: it's not the case that 'p' is true only if it's determinate that p. But this could be considered a high cost. Specific versions of onticism—like Barnes' view (2010)—might be able to dodge this bullet, but a defense of particular versions of onticist moral realism is beyond the scope of this article. In general, it seems that, in light of arguments like the one given just now, moral realists would have much work to do to show that moral realism coheres well with onticism. Moreover, while recently gaining followers, onticist theories are hardly the most popular theories of vagueness out there. Moral realists might thus better seek refuge somewhere else and defer to epistemicist or semanticist theories of vagueness.

Epistemicists usually believe that vague predicates have sharp boundaries and don't have a 'penumbral' extension: a man with 1856 hairs is either bald or not; Cheryl is either permitted to abort or not. For epistemicists, predicate vagueness gives rise to a particular sort of inescapable ignorance. Williamson (1996) explains this ignorance by appealing to shifty semantics: minor shifts in the usage of the vague predicate would change which property it refers to. Then, while Cheryl is either permitted to abort or not, it's inescapably epistemically risky to believe either the one or the other. Note that—and here's the crux Schoenfield (2016) overlooks in her critique of epistemicism—Williamson holds that one could not gain knowledge about borderline cases from data about linguistic usage: "Although meaning supervenes on use, there is no algorithm for calculating the former from the latter" (1996, p. 206).

One might account for vagueness-related ignorance by appealing to indeterminacy. If it's indeterminate whether Cheryl is permitted to abort, we could never know whether Cheryl is permitted to abort or not. But this move seems unavailable to epistemicist moral

---

<sup>41</sup> Note that anyone who claims that moral vagueness gives rise to indeterminacy (and that includes many semanticists) or that morality has other pockets of indeterminacy faces the same problem. One way out here is to bite the bullet and accept that there can be *truth-value gluts* or *true contradictions* and endorse *dialetheism* (Priest, 1987). However, as a matter of informal sociology of philosophy, a rejection of dialetheism seems to pervade the profession.

realists, who would typically hold that there *is* a fact of the matter of whether Cheryl is permitted to abort.<sup>42</sup>

However, something seems amiss about the view that we can be inescapably ignorant about what's permissible to do if there's a fact about what is permissible to do. Sorensen remarks that having unknowable obligations seems unfair (1991, p. 306). Moreover, many contend that moral facts are intrinsically action-guiding, motivating, or reason-giving, which might seem at odds with there being unknowable ethical facts.<sup>43</sup>

In addition, epistemicist moral realists face the following problem: If moral properties had sharp boundaries, a tiny difference in a property the moral property supervenes on could make a vast moral difference. Until a specific time, Cheryl could still abort, but not one nanosecond later. But shouldn't the quality of changes in a supervenience relation be proportional?<sup>44</sup> In other words, moral epistemicism entails that moral properties are *hypersensitive*, as Pummer (2022) puts it. Property *A* is hypersensitive iff *A* supervenes on property *B*, and a slight difference in *B* makes a significant difference in *A* by itself. Some authors don't think hypersensitivity is odd: micro-physical differences collapse buildings (Kultgen, 2018, p. 65), and a bit more sparkle can make a kitchen look fantastic (Hawthorne, 2022, p. 243). However, Kultgen seems to confuse hypersensitivity with a similar phenomenon, where small changes in B-properties make significant differences in A-properties by causing bigger differences in B-properties (see Pummer, 2022, p. 512), and intuitions might differ regarding Hawthorne's case. Hypersensitivity remains problematic.

In sum, moral realists are not clearly justified in passing on moral vagueness to epistemicist theories. At worst, this strategy might amount to jumping out of the frying pan and into the fire. Moral realists would then have to handle two new problems: Namely, it seems that there are no unknowable moral facts and that moral properties are not hypersensitive. Let's then turn to semanticism.

Semanticism is the view that vague predicates (of our actual language) lack precise application conditions. In a perfect language, all predicates would have semantic rules

---

<sup>42</sup> Epistemicists cannot appeal to indeterminacy if we understand under indeterminacy "some kind of status other than truth and falsity—a lack of truth-value, perhaps, or the possession of some other truth-value" (Wright, 2001, p. 69)).

<sup>43</sup> See, for instance, Constantinescu (2014).

<sup>44</sup> See Constantinescu (2014, p. 182); Dougherty (2014, p. 361).

that specify which properties they refer to. Some semanticist theories appeal to stance-dependent vague properties and thus appear to be non-starters for moral realists.<sup>45</sup>

Supervaluationism is arguably the most popular theory of vagueness. So, it should be the starting point for semanticist, moral realists. Supervaluationists claim that vague predicates are deficient in meaning due to having several admissible precisifications—i.e., different ways of making the predicate precise are appropriate.<sup>46</sup> Because different precisifications pick out different properties, the semantic rules of our language do not fix which property vague predicates like ‘... is bald’ refer to. Instead, there are multiple acceptable candidates for sets of instances to which one can attribute such predicates, and borderline cases are those cases that are in some but not all of those sets. Somewhat colloquially: we haven’t made up our minds on whether a man with 1856 hairs counts as bald and whether abortion at 151 days counts as permissible. Or so supervaluationists would argue.

The most pressing problem for supervaluationist moral realists is as follows. Suppose Cheryl deliberates on the semantics of ‘... is permissible’ and learns all reference-fixing facts about ‘... is permissible’. According to supervaluationism, shouldn’t she then come to know whether she’s permitted to abort, not allowed to abort, or whether it’s borderline whether she can abort? Schoenfield gives a similar (sometimes misunderstood<sup>47</sup>) example and concludes that the above picture is a bizarre view of moral deliberation. In her example, an agent arrives at moral knowledge by ‘crunching through the data’ about the use of ‘... is permitted’ in her community (2016, p. 265). Kultgen (2018) and Abasnezhad (2019) rightfully argue that semanticists needn’t commit to the view that such data would suffice to determine whether ‘... is permissible’ applies to a particular situation. However, moral supervaluationism and other forms of moral semanticism seem liable to the view that linguistic knowledge brings about moral knowledge. And this still appears to conflict with moral realism insofar as this position entails that an investigation of the meanings and the linguistic conventions of ethical vocabulary cannot reveal the truth value of moral judgments.

---

<sup>45</sup> See, for instance, Shapiro’s (2006) approach, which Constantinescu (2014, pp. 162-6) discusses.

<sup>46</sup> See Fine (1975). Sud (2019) offers an alternative semanticist approach for stance-independent realists.

<sup>47</sup> For instance, Abasnezhad (2019) overlooks that the person in Schoenfield’s example discovers it’s *not* borderline whether she can abort.

So, moral realists are not clearly justified in deferring moral vagueness to semanticist theories either. Again, this strategy would involve the necessity to handle a new problem: Namely, it seems that moral epistemology is independent of linguistics. In section 5.3, I will discuss another problem that moral realists face when passing on the datum of moral vagueness to supervaluationists. That problem can, however, also be seen as a data-disabler. So, I'll reserve discussion of it for later.

We have seen that it's not unproblematic for moral realists to pass the datum that there is moral vagueness on to philosophers who work on vagueness. And the problems raised are undoubtedly just the tip of a sizeable iceberg. Fortunately, this thesis is not concerned with solving any of these problems. The objective of this section was merely to showcase some of the issues that taking moral vagueness seriously brings for metaethical theorizing: moral vagueness is a datum that cannot easily be passed on. This is the first way in which moral vagueness is hard to handle.

## 5.2 *Data-Disablers of the First Kind*

In his (2018), Kultgen proposes a precautionary methodological stance that he calls "Vagueness Last." For Kultgen, we must explore all other plausible explanations before diagnosing a predicate with vagueness. 'Vagueness Last' is supposed to ensure that we don't hastily concede to vagueness without thoroughly considering alternative possibilities. Why such caution? According to Kultgen, "arguments from vagueness are ... easy to generate, and their philosophic power, often destructive, is impressive." (p. 68). Here's one of his examples of such arguments:

**P1** If moral realism is true, then '... is permissible' must be precise.

**P2** '... is permissible' is not precise.

Therefore, moral realism is false.

To avoid such 'impressively powerful and destructive' arguments, Kultgen advises a thorough examination of how things seem to us before we diagnose a predicate with vagueness. Not being hasty in one's phenomenological investigation is obviously important. Hasty phenomenologists produce unreliable seemings. So, Kultgen warns us that the way the datum that there's moral vagueness was produced (via seemings) might

have been faulty—that is, he aims to present what B, C & S-L would call data-disablers of the first kind. My aim here is to argue that Kultgen’s data-disablers are not decisive.

Before discussing Kultgen’s (apparent) data-disablers, let me briefly mention two ways to understand his critique that would *not* count as data-disablers. First, data are not disabled simply because they would have destructive implications. One shouldn’t ignore a cancer diagnosis just because it would be bad to have cancer. A theory must handle the data, not simply dismiss them on the grounds that it can’t handle them. Second, suppose one were to understand Kultgen as saying that one should be cautious in diagnosing moral predicates with vagueness because this would *inevitably* have destructive consequences. If that’s what we would take away from Kultgen, we would neither have identified a data-disabler, simply because no philosophical datum is inevitably destructive. Every philosophical domain has data that speak against otherwise very strong theories. This, however, just means that those theories must become even stronger by handling such data.

Kultgen presents a list of ‘tests’ one should run before diagnosing a predicate with vagueness: The Kripke-Test, the Computer-Test, the Parasitism-Test, the Moore-Test, the Soriticality-Test, and the Disagreement-Test (see pp. 67ff.). I will only briefly discuss the latter two now and go on to discuss the remaining four in detail after that.

First, the Soriticality-Test. Kultgen argues that many cases only appear soritical, but on closer inspection, they turn out to be merely *prima facie* soritical. It will be clear that the Soriticality-Test does not provide a strong data-disabler. I argued above that the datum that there’s moral vagueness can be derived without appealing to soriticality. So, pointing out mere *prima facie* soriticality has no bearing on how the datum was produced. Moreover, soriticality is not even a clear indicator of vagueness, as I argued in Chapter 3.

Next, the Disagreement-Test. Kultgen argues that in cases where we know that a case is a borderline case, we would have “nothing to talk about” (p. 67). In alleged moral borderline cases, however, we still wonder and disagree about whether the term applies. Our seeming that there’s moral vagueness might have been too hasty as we haven’t taken possible disagreement into account. Here’s a short answer to this worry: Kultgen confuses underdetermination-related borderline cases with vagueness-related borderline cases. In the former, we indeed have nothing to talk about. In the latter, we feel torn: we face conflicting reasons when we must categorize vagueness-related borderline cases. To repeat, we “have some inclination to judge it to have the property, and some inclination to judge it not to have the property” (Schiffer, 2003, p. 5). I’ll further elaborate on this point in the following section.

Let's now go through the remaining four tests, one by one, to see whether they provide us with strong data-disablers.

### 5.2.1 *The Kripke-Test*

What Kultgen calls the Kripke-Test boils down to the following: If a predicate seems vague to you, imagine a world in which the predicate in question is *not* vague.<sup>48</sup> Suppose the predicate would still seem vague in such a world. Then, Kultgen contends, our seeming that a predicate is vague does not provide evidence that this predicate is vague *in our world*. If a seeming that  $p$  is fully compatible with  $p$  and  $\sim p$ , it doesn't produce the datum  $p$  (or so Kultgen would say). And does the predicate '*... is permissible*' still seem vague in a world where it's not vague? "Yes, it does", says Kultgen (p. 55).

I don't believe the Kripke-Test provides a strong data-disabler. Note that Kultgen does not distinguish between the phenomenology of vagueness and that of ordinary uncertainty. He holds that non-vague predicates for which we are in the dark about *what* applies to them also appear vague. That might be true in some cases. However, the phenomenology of ordinary uncertainty is usually different. And taking that into account, the Kripke-Test loses much of its initial appeal.

To illustrate, let's follow Kultgen and try to imagine a world where '*... is permissible*' is not vague. In some cases of ordinary uncertainty about what a particular predicate applies to, the predicate seems well-functioning. It seems that it has a function mapping subjects to either its extension or anti-extension. This is clearly the case when we know a predicate's function (for instance, that of '*... is prime*'). If it seems to the inhabitants of the world Kultgen wants us to imagine that '*... is permissible*' is well-functioning, it will not seem vague to them. In other cases of ordinary uncertainty, it might not be clear that the predicate in question is well-functioning. However, if it nevertheless seems to the inhabitants of the world Kultgen wants us to imagine that something could

---

<sup>48</sup> Kultgen calls it Kripke-Test because of the following passage from Kripke:

I propose the following test for any alleged counterexample to a linguistic proposal: If someone alleges that a certain linguistic phenomenon in English is a counterexample to a given analysis, consider a hypothetical language which (as much as possible) is like English except that the analysis is stipulated to be correct. Imagine such a hypothetical language introduced into a community and spoken by it. If the phenomenon in question would still arise in a community that spoke such a hypothetical language . . . then the fact that it arises in English cannot disprove the hypothesis that the analysis is correct for English. (1977, p. 265)

be forthcoming to resolve whether ‘... is permissible’ applies to cases such as Cheryl’s or Darryl’s, ‘... is permissible’ also won’t seem vague.

So, what would be the phenomenology of the inhabitants of a world where ‘... is permissible’ is not vague? Would ‘... is permissible’ feel well-functioning? Would cases like Cheryl’s or Darryl’s feel resolvable? Or would ‘... is permissible’ still feel vague? I’m not sure how to answer these questions. In fact, I’m not sure there’s a way to answer these questions that is not utterly speculative. Maybe even after the phenomenological clarification I offered, the inhabitants of that world would still feel that ‘... is permissible’ is vague. Maybe not.

A more direct way of making essentially the same point is this. What counts is whether we in *our* world have a distinct seeming that moral predicates are vague. Contrary to Kultgen, I believe there is a phenomenological contrast between vagueness and ordinary uncertainty. Moreover, I believe that when we become clearer about that contrast, many philosophers will indeed report that it seems to them that moral predicates are vague. Of course, just because a predicate seems vague doesn’t mean it is. However, the Kripke-Test does not provide a strong reason to believe that we systematically confuse ordinary ignorance with vagueness.

### 5.2.2 *The Computer-Test*

Here is how Kultgen lays out the Computer-Test.

When I was young, I badly wanted the Enterprise’s computer from Star Trek: The Next Generation. All one had to do was say, “Computer” then ask a question and the computer would come up with the answer. Suppose there was a computer like that, only literally omniscient. Call it Computer. Computer can answer any question that has an answer. Importantly, Computer knows the correct moral theories, and can perfectly apply them. (p. 57)

The test is the following: Can we ask Computer anything to *indirectly* find out whether, say, Darryl is or is not morally permitted to look at his phone for 30 seconds? Kultgen holds that if yes, the case does not illustrate vagueness because vague predicates are *inquiry-resistant*. We could, for instance, not find out whether a man with 1856 hairs

on his head is hirsute by asking Computer questions. Can we learn something from Computer about Darryl's case? According to Kultgen, yes:

We could learn what the true moral theory is. We could learn the personal histories of every individual within 10 miles of the park while Darryl and his child are there. We could learn that there is a stealth kidnapper at the park who will steal Darryl's child if he looks away for any longer than 7 seconds. (p. 59)

This information, Kultgen contends, could help us to know whether Darryl is or is not permitted to look at his phone for 30 seconds.

I think the Computer-Test also fails to provide a strong data-disabler. First, Kultgen does not take Williamsonian epistemicists into account. According to Williamsonian epistemicists, having a good talk with an omniscient being *can* provide us with knowledge about borderline cases. And maybe Williamsonian epistemicism is the correct view on vagueness. So, the Computer-Test doesn't tell us whether a predicate is vague.<sup>49</sup>

Second, and maybe more importantly, if Darryl knew all the things we could learn from Computer, he would simply be in a different situation. Consider the following case:

Darryl is permitted to look at his phone for 30 seconds while knowing that there's a stealth kidnapper at the park who will steal Darryl's child if he looks away for any longer than 7 seconds.

Of course, it's not borderline whether Darryl is permitted to look at his phone for 30 seconds in *that* situation. He's not. But the seeming that '... is permissible' is vague was not produced by reflecting on that situation. It involved reflecting on the original situation

---

<sup>49</sup> Kultgen could reply that Williamsonian epistemicists are not committed to the idea that we could learn whether a borderline case applies to a vague predicate via indirect information but only by *directly* asking an omniscient being whether the case applies. Recall that Williamson, for instance, holds that one could not gain knowledge about borderline cases from information about linguistic usage (see 5.1). That granted, Williamsonians might still say that an omniscient being (the only being that can *know* whether a borderline case applies, in Williamson's view) could be able to give us indirect enough information so that we could justifiably say that we have found out what the case is. However, the precise role of omniscient beings in Williamson's view is a difficult interpretive question that I will not try to address here.

where Darryl is not informed by an omniscient being. So, one might say the Computer Test is irrelevant.

But maybe that's too fast. Kultgen could argue that what we learned from Computer shows that even in his original, ignorant situation, Darryl is not permitted to look at his phone for 30 seconds. What Darryl knows does not matter, and neither what he believes nor what other internal features he has. It is the *fact* that there's a kidnapper in the park that makes looking at his phone for 30 seconds impermissible. And that is so, Kultgen could claim, whether Darryl is aware of that fact or not.

Here are two points addressing this reply. First, one might grant that what *we* know is, unbeknownst to Darryl, a decisive reason for him not to look at his phone for 30 seconds. It wouldn't be *irrational* for Darryl to look at his phone for 30 seconds, however. What's rational to do depends on what we believe.<sup>50</sup> One could then continue by insisting that if an action is not irrational, that action can't be morally prohibited. This view of the relationship between rational and moral action is not obviously the right one. But I submit that it is not obviously the wrong one either.

Second, one might complain that the idea that we could learn from Computer facts, which entail that it's impermissible for Darryl to look at his phone for 30 seconds, is unsubstantial. If a case seems morally vague, it seems the matter can't be resolved if all relevant facts are known. Conversely: if it seems to Kultgen that we could learn from Computer facts that could resolve the question, he merely appears to say that it doesn't seem to him that Darryl's case illustrates moral vagueness. In other words, Kultgen's appeal to the Computer Test only reports an opposing seeming in masquerade.

In sum, the Computer Test, I believe, also does not provide strong reasons to believe that the datum that there's moral vagueness was produced might have been faulty.

### 5.2.3 *The Parasitism-Test*

The Parasitism-Test rests on the idea that *every* predicate can have borderline predications. One just has to insert a *borderline subject*. I have already mentioned Constantinescu's and Kultgen's examples above.

This cloud weighs one billion kilograms.

---

<sup>50</sup> I here follow Parfit (2011, p. 34f.)

The number of heaps in the sandbox is prime. (Imagine four heaps and one borderline heap in the sandbox)

Both examples seem like borderline predications, but neither ‘... weighs one billion kilograms’ nor ‘... is prime’ are vague predicates. Are apparent moral borderline predicates merely parasitic on borderline subjects?

There are two ways to understand the Parasitism-Test. Maybe the test is supposed to show that there’s no *genuine* moral vagueness. Moral vagueness just arises from ordinary vagueness. Here, I’m inclined to answer with Sorensen (1991): “But so what? Parasitic borderline cases are still borderline cases” (p. 307). Or maybe the Parasitism-Test is supposed to show that borderlineness is not sufficient for vagueness. As should be clear, I agree. In any case, the Parasitism-Test does not show that the seeming that there’s moral vagueness was hastily produced. ‘The number of heaps in the sandbox is prime’ is clearly phenomenologically different from ‘Darryl is permitted to look at his phone for 30 seconds’. Reflecting on the sentence ‘The number of heaps in the sandbox is prime’ does not highlight that something feels odd about the function of ‘... is prime’. But reflecting on ‘Darryl is permitted to look at his phone for 30 seconds’ does highlight that ‘... is morally permitted’ is odd, however.

#### 5.2.4 *The Moore-Test*

Finally, the Moore-Test:

According to the Moore Test, if the vagueness of some predicate “F” is in question, formulate two arguments; one demonstrating the most implausible consequence of “F” being precise, and the other demonstrating the most implausible consequence of “F” being vague. Then, evaluate whether it is more implausible that “F” is vague or more implausible that it is precise. (p. 66).

Here is what Kultgen believes those arguments are:

**P1** If ‘... is permissible’ is not vague then there is a sharp cutoff where it becomes impermissible for Darryl to look at his phone.

**P2** ‘... is permissible’ is not vague.

Therefore, there is a sharp cutoff where it becomes impermissible for Darryl to look at his phone.

**P1** If ‘... is permissible’ is vague then moral realism, moral naturalism, utilitarianism, Kantianism, expressivism, and probably more, are clearly false.

**P2** ‘... is permissible’ is vague.

Therefore, moral realism, moral naturalism, utilitarianism, Kantianism, expressivism, and probably more, are clearly false.

For Kultgen, it’s more implausible that moral realism, moral naturalism, utilitarianism, Kantianism, expressivism, and probably more, are clearly false. So, he claims that we should side with the view that ‘... is permissible’ is not vague.

I don’t agree that either of the above conclusions can be drawn so easily. Moreover, as I said at the beginning of this section, data are not disabled simply because they would have destructive implications. It’s worth, however, to shortly discuss an idea that might be in the background of Kultgen’s thought. This idea is that the theories that are threatened by moral vagueness are simply too good to give up because of some annoying seeming.

Imagine a theory *T* that—before the datum that there’s moral vagueness is collected—does everything well according to the criteria of B, C & S-L. It accommodates and explains all the collected data thus far. It also excels on the levels of theorizing that follow the initial level of handling the data. It substantiates and integrates the data, for instance, by offering positive epistemic support for its claims, responding to objections, and being coherent. Moreover, *T* possesses specific theoretical virtues such as simplicity, modesty, and beauty. Now imagine moral vagueness is a consideration that *T* can’t handle. Can’t such a great theory be excused to ignore one datum?

If your answer is yes, you’re not on board with the framework of B, C & S-L. For B, C & S-L, there’s no escaping the data. What was a great theory until the discovery of a new datum must handle that datum as much as any other theory. The idea that *T* can be excused from handling the consideration that there’s moral vagueness presupposes a different framework. For instance, advocates of the cost-benefit method might support this idea. Here’s how B, C & S-L define the cost-benefit method:

When constructing a theory about a given domain, theorists ought to identify a set of theses about the domain that possess a high ratio of positive features (benefits) to negative ones (costs), where these features are in themselves roughly equal and aggregative. The best theory is the one that achieves the highest such ratio relative to rivals. (p. 93).

For the reasons given above, our theory *T* has a high number of what Enoch (2011, p. 267) calls “plausibility points.” If a competing theory can handle moral vagueness, it would gain some points. But as long as *T* still has the most points, it will continue to be the greatest.

In the framework of B, C & S-L, Kultgen’s Moore-Test is insignificant. I defended B, C & S-L’s framework (or rather my use of it) in 4.1.4. But let me use this occasion to express explicitly that I think that the framework of B, C & S-L is preferable to the cost-benefit method. Something seems amiss about distributing plausibility points for fulfilling very different theoretical tasks. As B, C & S-L express it, the cost-benefit method seems to allow for “illicit trade-offs” (p. 98). As one example, the cost-benefit method seems to allow for dismissing annoying data by making the theory more beautiful. Certainly, to again use the terminology of B, C & S-L, a sound philosophical method must be *hierarchical*. It must recognize “that criteria play diverse roles, and that some roles possess more significance than others” (p. 89).

### 5.3 *A Data-Disabler of the Second Kind*

In the preceding section, I argued that Kultgen’s ‘Vagueness Last’ tests do not provide strong data-disablers of the first kind. This was the second way in which moral vagueness is hard to handle.

In this last section, I’m going to argue against a data-disabler of the second kind. Thus, moral vagueness is also hard to handle in that way. Specifically, I’m going to argue for something that some will find odd but which I can’t help thinking is true: Namely, that, against the view of eminent philosophers such as Ruth Chang, borderline cases cannot be stipulated away.

Imagine you must categorize people as ‘bald’ or as ‘not-bald’ based only on the *metaphysical question* of whether they’re really bald. No extrinsic reasons (rewards, negative consequences, peer pressure, etc.) bear on your categorization choices. For each

person you encounter, you have extensive knowledge of facts relevant to whether they're bald (for instance, you know the number, length, and distribution of hairs on each person's head), and you're competent in using the predicate '... is bald.' Here comes Herbert, a borderline bald man. What now?

To many, it will seem that in such a situation, you can simply stipulate that Herbert is bald or that he's not. The consideration that we can stipulate borderline cases away is, in our context, a data-disabler of the second type. To see why, ask yourself whether Cheryl or Darryl could simply stipulate that '... is permissible' applies (or doesn't apply) in their situation. It seems ludicrous to claim that they could. So, we identified a competing consideration that potentially disables the datum that there's moral vagueness. In this section, I argue that it's not the case that we can stipulate borderline cases away—that is, I'm trying to disable this potential disabler.

The view that we can stipulate borderline cases away is most adamantly defended by Ruth Chang. According to Chang, the correct answer to whether Herbert is bald is that it's indeterminate whether he is (i.e., it's neither true nor false that he is). Yet she claims that “resolution of this indeterminacy can be appropriately given by a coin flip” (2001, p. 138). Importantly, Chang means by this not only that you can decide what to do with Herbert by fiat but that you can *close the metaphysical question* of whether Herbert really is bald through arbitrary stipulation (2022a, p. 55).

One obvious immediate response here is that epistemicists will not share Chang's view. That is simply for the reason that epistemicists don't believe that borderline cases are cases of indeterminacy. Let's look past this, however. This section discusses the intuition that borderline cases can be stipulated away. This intuition speaks against moral vagueness *and* against epistemicism.

As I'll explain in a moment, supervenience can readily explain the intuition that borderline cases can be stipulated away. In fact, Chang seems to believe that it *follows* from supervenience that we can stipulate what the case is in situations like the one above. One could then argue that moral theories cannot pass on the datum that there's moral vagueness to supervenience. That's an additional reason against combining moral theories with supervenience.

Here, however, I'm interested in framing a data-disabler of the second kind. One way to do so is simply to appeal to the seeming that one can stipulate borderline cases away. However, at first, it might seem that one can do better. Supervenience can explain why we could stipulate borderline cases away, and supervenience is a (maybe still

*the*) leading theory on vagueness. So, here's a sophisticated data-disabler consisting of the following set of considerations:

{Borderline cases can be stipulated away; One of our best theories of vagueness readily explains why that's the case}.

To find a way to tackle this disabler, let's first think about what it might mean to close a metaphysical question. Intuitively, one can close a metaphysical question by knowingly providing the correct answer and by knowingly responding that the question makes no sense. If you know it's raining, you can close the question of whether it's really raining by answering that it's raining. If you know that the question of whether fried eggs sound better than the number nine is absurd, you can close the question by rejecting it as nonsense.<sup>51</sup>

Assuming with Chang that the correct answer to whether Herbert is bald is that it's indeterminate whether he is, you cannot provide the correct answer in the situation described above. Nor can you reject the question (you must answer). But, according to Chang, there's a third option. You can close the metaphysical question of whether Herbert is bald by stipulating that he's bald or by stipulating that he's not.

As said, many will find this idea intuitive. But additionally, it seems that supervaluationism, one of the best theories of vagueness, can explain why we could flip a coin on whether Herbert is bald. Let's remember what story supervaluationists tell. The predicate '... is bald' expresses the concept *Baldness*. According to supervaluationism, *Baldness* has imprecise application conditions. There are some instances—like Herbert—where it's indeterminate whether *Baldness* applies. However, according to the supervaluationist, '... is bald' also has multiple admissible precisifications that can *sharpen* it. One precisification of '... is bald' might express *Baldness*<sub>1</sub>, which refers to the property *BALDNESS*<sub>1</sub>. Everyone with less than 1000 hairs on their head is bald<sub>1</sub>, and everyone with more than 1000 hairs is not. Another might express *Baldness*<sub>2</sub>, which refers to the property *BALDNESS*<sub>2</sub>. The extension of this precisification includes everyone with less than 1001 hairs on their head. Another precisification might express *Baldness*<sub>3</sub>, where

---

<sup>51</sup> To my knowledge, there's no standard view on question-closing. For a recent account of question-rejection, see Millson (2021).

everyone with less than 1002 hairs on their head is bald<sub>3</sub>. Let's say Herbert has 1001 hairs on his head. It's indeterminate whether he is bald; he's neither bald<sub>1</sub> nor bald<sub>2</sub>; he is bald<sub>3</sub>.

From a supervenience perspective, when we say that the predicate '... is bald' applies to Herbert, we might mean (a) that Herbert is in the extension of all admissible precisifications of '... is bald' or (b) that Herbert is in the extension of at least one admissible precisification of '... is bald'. To borrow terminology from Egré et al. (2013): Herbert isn't *strictly* bald (as (a) is false), but he's *tolerantly* bald (as (b) is true).<sup>52</sup> The supervenience explanation of the stipulation-intuition is that because Herbert is tolerantly bald and tolerantly not-bald, we can arbitrarily stipulate whether he's bald and thereby close the question of whether he is.

But what is meant by 'stipulation'? I'll discuss six candidate definitions of stipulation. I'll argue that only the sixth definition seems initially promising for Chang. Let the definiendum be *S stipulates that a is F*.

i. =<sub>df</sub> *S plans to act as if a is F*

Planning to treat Herbert as if he's bald is a non-starter for Chang's purposes. At most, doing so only closes the *practical* question of what to do with Herbert; it leaves the *metaphysical* question unaddressed.

ii. =<sub>df</sub> *S conditionally defines a as being F*

By 'conditionally defining,' I mean 'defining for the purposes of ...'. Such definitions do not close unconditional, metaphysical questions. So, conditionally defining Herbert to be bald is another non-starter.

iii. =<sub>df</sub> *S introduces the name 'a' for every x that is F*.

---

<sup>52</sup> This use of 'tolerant' is not to be confused with tolerance as an (alleged) feature of vague predicates. There are other ways of saying essentially the same thing as "Herbert is tolerantly bald" in this context: it's not true that Herbert is determinately bald, but it's true that he's not determinately not-bald; it's not super-true, but true that he's bald.

Sometimes, one might be able to close the question of whether  $a$  is  $F$  by introducing ‘ $a$ ’ as referring to every  $x$  that is  $F$ . (Did Jack the Ripper murder her? Yes, let’s call her murderer ‘Jack the Ripper.’) For borderline cases such as Herbert, however, it seems there’s nothing to introduce.

iv. =<sub>df</sub>  $S$  declares  $a$  to be  $F$

By ‘declaring,’ I mean an illocutionary act that, if performed successfully, “brings about the correspondence between the propositional content and reality” (Searle, 1975, pp. 16f.). Declaring can only be successful if you have the necessary authority. Imagine I’m in charge of the rules of a ball game and that the rules established so far do not determine whether a ball that lands straight on the goal line is a goal. By my authority, I declare that it’s a goal. The metaphysical question of whether it’s a goal seems closed.

Does Chang think one can bring about the correspondence between the propositional content of ‘Herbert is bald’ and reality? Apparently, no. Whether Herbert is bald, Chang claims, cannot be determined via coin flip: “The answer to that question is that it is indeterminate whether he is” (2002, p. 138). At least if one excludes the possibility of linguistic dictators with absolute authority to precisify vague words at will, it seems correct to reject declaration as a way of closing the question on borderline cases.

v. =<sub>df</sub>  $S$  declares that the question of whether  $a$  is  $F$  is the question of whether  $a$  is  $F_x$  (where ‘ $F_x$ ’ is a non-vague predicate)

There are two worries regarding this option. First, suppose I declare that the question is ‘Is Herbert bald<sub>3</sub>?’ and answer positively. Then *that* question has been closed—not the metaphysical question of whether Herbert is bald. Moreover, it hasn’t been closed by arbitrary stipulation but by knowingly giving the correct answer.

Second, if I declare that the question of whether Herbert is bald is ‘Is Herbert bald<sub>3</sub>?’ and you declare it is ‘Is Herbert bald<sub>2</sub>’, we seem to be in genuine conflict—yet Chang says that it could hardly be said that we disagree when we stipulate in opposing ways (2002, p. 138). Like Kultgen, she believes there’s nothing to disagree about in borderline cases. That’s, in fact, one of her motivations for the view that we can stipulate borderline cases away. Since opposing declarations seem to entail conflict, the idea that stipulating means declaring what the question is does not work well for Chang.

vi.  $=_{df}$  *S freely decides that a is F*

What I mean by ‘freely deciding’ here is best shown by example. Again, imagine a ball game where the rules do not determine whether a ball that lands straight on the goal line is a goal. This time, the rules are settled. If we subtract all extrinsic reasons to decide one way or the other, it seems we’re left with nothing to base our decision on. I believe that in such a case, the metaphysical question of whether *a* is *F* can indeed be closed via coin flip. There are simply no reasons to do anything else. We’re free to decide in the sense that we have no reasons for or against deciding in favor of one available option in particular. This example illustrates what could be called ‘stipulation as a free decision,’ and moreover, a situation in which stipulation, defined in this way, can close a metaphysical question on a case of indeterminacy.

At least initially, it seems that this definition might serve Chang’s purposes. But there’s a problem. We cannot make *arbitrary* decisions about *vagueness*-related borderline cases. In the ball game case, we lack intrinsic reasons to decide one way or the other, but we seem to have an intrinsic reason to categorize Herbert as bald. At the same time, we do have an intrinsic reason to categorize him as not-bald. As I laid out in section 4.4, we are typically resemblance-ambivalent about borderline cases. On the one hand, Herbert is *baldish*: he seems bald; he’s similar to clear cases of baldness; he’s sort of bald. On the other hand, he’s not-baldish. The ball game case, however, is not such. Instead, it is better characterized as an underdetermination-related borderline case.

One way to answer on behalf of proponents of the Changian data-disabler is to argue that resemblance-ambivalence does not provide real reasons to categorise Herbert as bald and as not bald after stipulation. Indeed, even if we’re typically ambivalent when confronted with a borderline case, it’s not necessarily fitting to still feel that way after stipulation. There can be a *psychological remainder*, but—Chang implies—it might rather be fitting to be a “confident stipulator” (2022a, p. 55). Put otherwise, one might claim that resemblance-ambivalence can (and maybe should) be overcome through stipulation. If that were the case for real borderline cases, we’d again have reasons to doubt that there’s moral vagueness.

Remember that I gave a conceptual and an empirical reason to believe that we’re typically resemblance-ambivalent about borderline cases in 4.4. The conceptual reason was that the gradability of borderline cases speaks in favor of the view that we’re

resemblance-ambivalent about borderline cases. The problem with Chang's view is that it's plausible that a borderline case remains gradable in relation to other cases after stipulation. After I stipulated that Herbert is bald, he'll still be to a degree similar to clear cases of baldness and to clear cases of baldness. So, we would still have intrinsic reasons to answer whether Herbert is bald both positively and negatively after stipulation.

It seems Chang now would have to argue that resemblance-ambivalence doesn't provide reasons to categorize Herbert as bald and as not bald even *before* stipulation.

Note that supervaluationists can account for degrees of borderlineness by appealing to what may be called a *sharpening ratio*. To illustrate the idea, suppose the only admissible sharpenings of '... is bald' express *Baldness*<sub>1-3</sub>. Herbert has 1001 hairs on his head, so he's neither bald<sub>1</sub> nor bald<sub>2</sub>, but bald<sub>3</sub>. So, the sharpening ratio in that case would be 1/3. Now, if the sharpening ratio works out that way, then supervaluationists could hold that Herbert is, in a relevant sense, more not-bald than bald. But, Chang writes, "appealing to a ratio of possible sharpenings to resolve the indeterminacy is to appeal to something extrinsic to the facts about how Herbert stands to the word 'bald', such as 'majority rule'" (2022b, p. 403). Chang here seems to claim that if Herbert is 1/3 bald according to the sharpening ratio, we nevertheless don't face a situation where we have slightly more, but not decisive, reason to categorize Herbert as not-bald than as bald. The sharpening ratio does not mean that Herbert is more similar to not-bald than to bald men. Whether in the minority or not, each admissible sharpening is equally acceptable. So, Chang might have a story about why we might think that we have competing categorization reasons when facing borderline cases. We confuse sharpening-ratios with similarity.

However, now recall the empirical reasons I gave above for the view that we're resemblance-ambivalent about borderline cases. The above-cited empirical evidence gives us robust reasons to believe that most people treat borderline cases as cases of competing similarity. We have no empirical evidence that a larger or equal number of people feel that borderline cases can be stipulated away. So why think that we systematically confuse sharpening-ratios with similarity?

Moreover, note that the notion of a sharpening-ratio is not even supervaluationist standard vocabulary. Appealing to sharpening-ratios *could* explain the appearance of resemblance-ambivalence. However, supervaluationists don't have to follow Chang's claim that sharpening ratios do not correspond to categorization reasons. In light of the empirical evidence, they might be better advised to embrace resemblance-ambivalence

and claim that sharpening-ratios do correspond to degrees of similarity. In other words, supervaluationists do *not* have to support Chang's view that borderline cases can be stipulated away. We thus arrived at a point where the sophisticated data-disabler consisting of the set of considerations {Borderline cases can be stipulated away; One of our best theories of vagueness readily explains why that's the case} falls apart. And, in light of the evidence that most people do treat borderline cases as cases of competing similarity, the less sophisticated data-disabler, consisting only of the consideration {Borderline cases can be stipulated away} seems not particularly strong.

## Chapter 6: CONCLUSION

Reconsider:

Cheryl considers an on-demand abortion after 151 days of gestational age. Is Cheryl morally permitted to do so?

Darryl is watching his toddler at the playground. Is he morally permitted to look at his phone for 30 seconds?

I believe that when reflecting on such cases, it will seem to many that the predicate '... is morally permitted' is vague. And similar examples can easily be created for other moral predicates. Thus, it will seem to many that there's moral vagueness. In this thesis, I hope to have given a convincing rationale for the idea that such seemings are enough to take the consideration that there's moral vagueness seriously.

I defended four claims:

*Claim 1:* There's no straightforward way to argue for moral vagueness.

*Claim 2:* Because it seems that moral predicates are vague, we have reason to believe that moral predicates are vague.

*Claim 3:* Vagueness has a distinct phenomenological profile.

*Claim 4:* Apparent reasons to dismiss moral vagueness despite the appearance that moral predicates are vague are unconvincing.

Regarding Claim 1, I first argued that straightforward arguments for moral vagueness presuppose that there's theoretical consensus of what vagueness consists in. That's, however, not the case: the points of disagreement between those who characterize vagueness as borderlineness, those who characterize vagueness as boundarylessness, and those who characterize vagueness via epistemic or psychological notions are deep and possibly intractable. I then showed that, on closer inspection, particular straightforward arguments are, for other reasons as well, not straightforward at all.

Regarding Claim 2, I argued that we can have a seeming that moral predicates are vague, even though we don't know what vagueness consists in. Insofar as it does, in fact, seem to many theorists that moral predicates are vague, 'There is moral vagueness' is a philosophical datum—a consideration that must be taken seriously.

Regarding Claim 3, I argued that vagueness can be phenomenologically characterized as giving rise to resemblance-ambivalence. I believe that this phenomenological clarification can help people to realize what it's like to have the seeming that a predicate is vague. I also believe that when reflecting on the what-it's-likeness of vagueness, more people will conclude that it seems that there's moral vagueness.

Regarding Claim 4, I argued that moral realists cannot simply pass on the datum of moral vagueness to philosophers of vagueness. Neither onticism, nor epistemicism, nor semanticism clearly coheres with moral realism. So, moral realists cannot 'dismiss' moral vagueness in the sense of placing the burden of explanation on philosophers of vagueness. I also argued that none of Kultgen's 'Vagueness Last' tests provides a strong data disabler. While his tests might contribute to a more careful phenomenological judgment about cases like Cheryl's or Darryl's, they are also, in many aspects, misleading. Finally, I argued that the idea that borderline cases can be stipulated away neither provides a strong data-disabler. I believe that that idea rests on a failure to take into account that vagueness-related borderline cases are cases of resemblance-ambivalence.

This all is, of course, by no means the end of the debate on moral vagueness but rather merely its beginning. As I stated in the introduction, what started this thesis was the question, 'What should one do when it's vague what one should do?' I believe that this question is a question of utmost importance—for moral theory, as well as for all of us in our personal lives. Suppose it is indeed borderline permissible for Cheryl to have an abortion after 151 days and for Darryl to look at his phone for 30 seconds. But now what? What should we do when it's vague what we should do? This thesis was an attempt to put

these questions, which we feel pressing in our everyday lives and face in many of our hardest decisions, firmly on the agenda of moral theory.

## BIBLIOGRAPHY

- Abasnezhad, A. (2019). Moral Realism and Semantic Accounts of Moral Vagueness. *Inquiry*, 1–13. <https://doi.org/10.1080/0020174X.2019.1570866>
- Alxatib, S., & Pelletier, F. J. (2011). The Psychology of Vagueness: Borderline Cases and Contradictions. *Mind & Language*, 26(3), 287–326. <https://doi.org/10.1111/j.1468-0017.2011.01419.x>
- Aristotle. (1986). *De Anima (On the Soul)* (H. Lawson-Tancred, Trans.). Penguin Books.
- Armstrong, D. M. (1989). *Universals: An Opinionated Introduction*. Westview Press.
- Baima, N. (2014). The Problem of Ethical Vagueness for Expressivism. *Ethical Theory and Moral Practice*, 17(4), 593–605. <https://doi.org/10.1007/s10677-014-9493-0>
- Barnes, E. (2010). Ontic Vagueness: A Guide for the Perplexed. *Noûs*, 44(4), 601–627. <https://doi.org/10.1111/j.1468-0068.2010.00762.x>
- Barnett, D. (2009). Is Vagueness *Sui Generis*? *Australasian Journal of Philosophy*, 87(1), 5–34. <https://doi.org/10.1080/00048400802237376>
- Bengson, J., Cuneo, T., & Shafer-Landau, R. (2022). *Philosophical Methodology: From Data to Theory* (First edition). Oxford University Press.
- Bird, A. (2007). *Nature's Metaphysics: Laws and Properties*. Oxford University Press.
- Bogen, J., & Woodward, J. (1988). Saving the Phenomena. *The Philosophical Review*, 97(3), 303. <https://doi.org/10.2307/2185445>
- Broome, J. (1997). Is Incommensurability Vagueness? In R. Chang (Ed.), *Incommensurability, Incomparability, and Practical Reason*. Harvard University Press.
- Brown, P., & Levinson, S. C. (1987). *Politeness: Some Universals in Language Usage*. Cambridge University Press.
- Bueno, O., & Colyvan, M. (2012). Just What Is Vagueness? *Ratio*, 25(1), 19–33. <https://doi.org/10.1111/j.1467-9329.2011.00513.x>
- Burns, L. (1995). Something to Do with Vagueness. *Southern Journal of Philosophy*, 33(S1), 23–47. <https://doi.org/10.1111/j.2041-6962.1995.tb00761.x>
- Chang, R. (2001). *Making Comparisons Count*. Routledge.

- Chang, R. (2022a). Are Hard Cases Vague Cases? In H. Andersson & A. Herlitz (Eds.), *Value Incommensurability: Ethics, Risk and Decision-Making*. Routledge.
- Chang, R. (2022b). How to Avoid the Repugnant Conclusion. In J. McMahan, T. Campbell, J. Goodrich, & K. Ramakrishnan (Eds.), *Ethics and Existence* (pp. 389–429). Oxford University Press.  
<https://doi.org/10.1093/oso/9780192894250.003.0015>
- Conee, E. (1998). Seeing the Truth. *Philosophy and Phenomenological Research*, 58(4), 847–857.
- Constantinescu, C. (2014). Moral Vagueness. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Volume 9* (pp. 152–185). Oxford University Press.  
<https://doi.org/10.1093/acprof:oso/9780198709299.003.0007>
- Cucchi, C. (2007). An Investigation of General Extenders in a Corpus of Parliamentary Debates. In M. C. Paganoni & M. Bait (Eds.), *Discourse and Identity in Specialized Communication: Conference Proceedings: Palazzo Feltrinelli, Gargnano del Garda (Brescia), Italy June 25-26, 2007*. Lubrina.
- Deemter, K. van. (2010). *Not Exactly: In Praise of Vagueness*. Oxford University Press.
- DePaul, M. (2009). Phenomenal Conservatism and Self-Defeat. *Philosophy and Phenomenological Research*, 78(1), 205–212. <https://doi.org/10.1111/j.1933-1592.2008.00240.x>
- Descartes, R. (1993). *Meditations on First Philosophy: In Which the Existence of God and the Distinction of the Soul From the Body Are Demonstrated* (D. A. Cress, Trans.; 3rd ed). Hackett Pub. Co.
- Dougherty, T. (2014). Vague Value. *Philosophy and Phenomenological Research*, 89(2), 352–372. <https://doi.org/10.1111/phpr.12026>
- Dunaway, B. (2017). Ethical Vagueness and Practical Reasoning. *The Philosophical Quarterly*, 67(266), 38–60. <https://doi.org/10.1093/pq/pqw038>
- Earlenbaugh, J., & Molyneux, B. (2009). Intuitions Are Inclinations to Believe. *Philosophical Studies*, 145(1), 89–109.
- Edgington, D. (2016). Schiffer on Indeterminacy, Vagueness, and Conditionals. In G. Ostertag (Ed.), *Meanings and Other Things: Themes From the Work of Stephen Schiffer*. Oxford University Press.
- Egré, P., De Gardelle, V., & Ripley, D. (2013). Vagueness and Order Effects in Color Categorization. *Journal of Logic, Language and Information*, 22(4), 391–420.  
<https://doi.org/10.1007/s10849-013-9183-7>

- Eklund, M. (2005). What Vagueness Consists In. *Philosophical Studies*, 125(1), 27–60.
- Ekman, P. (1992). *Telling Lies: Clues to Deceit in the Marketplace, Politics, and Marriage*. W.W. Norton.
- Enoch, D. (2011). *Taking Morality Seriously*. Oxford University Press.  
<https://doi.org/10.1093/acprof:oso/9780199579969.001.0001>
- Fine, K. (1975). Vagueness, Truth and Logic. *Synthese*, 30(3–4), 265–300.  
<https://doi.org/10.1007/BF00485047>
- Geach, P. T. (1966). Plato’s Euthyphro: An Analysis and Commentary. *The Monist*, 50(3), 369–382.
- Greenough, P. (2003). Vagueness: A Minimal Theory. *Mind*, 112(446), 235–281.
- Grice, H. P. (1989). Logic and Conversation. In *Studies in the Way of Words* (pp. 22–40). Harvard University Press.
- Hacking, I. (1992). The Self-Vindication of the Laboratory Sciences. In A. Pickering (Ed.), *Science as Practice and Culture* (pp. 29–64). University of Chicago Press.
- Hampton, J. A., Estes, Z., & Simmons, C. L. (2005). Comparison and Contrast in Perceptual Categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(6), 1459–1476. <https://doi.org/10.1037/0278-7393.31.6.1459>
- Hawthorne, J. (2022). Moral Vagueness and Epistemicism. *Journal of Ethics and Social Philosophy*, 23(2). <https://doi.org/10.26556/jesp.v23i2.1602>
- Huemer, M. (2007). Compassionate Phenomenal Conservatism. *Philosophy and Phenomenological Research*, 74(1), 30–55. <https://doi.org/10.1111/j.1933-1592.2007.00002.x>
- Jucker, A. H., Smith, S. W., & Lüdge, T. (2003). Interactive Aspects of Vagueness in Conversation. *Journal of Pragmatics*, 35(12), 1737–1769.  
[https://doi.org/10.1016/S0378-2166\(02\)00188-1](https://doi.org/10.1016/S0378-2166(02)00188-1)
- Kalmus, H. (1979). Dependence of Colour Naming and Monochromator Setting on the Direction of Preceding Changes in Wavelength. *The British Journal of Physiological Optics*, 33(2), 1–9.
- Koester, A. (2007). ‘About Twelve Thousand or So’: Vagueness in North American and UK Offices. In J. Cutting (Ed.), *Vague Language Explored* (pp. 40–61). Palgrave Macmillan UK. [https://doi.org/10.1057/9780230627420\\_3](https://doi.org/10.1057/9780230627420_3)
- Kripke, S. (1977). Speaker’s Reference and Semantic Reference. *Midwest Studies in Philosophy*, 2(1), 255–276.

- Kultgen, B. M. (2018). *Vagueness and Ethics* [Harvard University].  
[https://scholar.colorado.edu/concern/graduate\\_thesis\\_or\\_dissertations/m613mx68w](https://scholar.colorado.edu/concern/graduate_thesis_or_dissertations/m613mx68w)
- Leonelli, S. (2016). *Data-Centric Biology: A Philosophical Study*. University of Chicago Press. <https://doi.org/10.7208/chicago/9780226416502.001.0001>
- Łukasiewicz, J. (2005). Many-Valued Systems of Propositional Logic. In S. McCall (Ed.), *Polish logic, 1920-1939*. Clarendon Press.
- Lycan, W. G. (1988). *Judgement and Justification* (Vol. 100, Issue 3, pp. 481–484). Cambridge University Press.
- Mácha, J. (2022). Hegel and Wittgenstein on Difficulties of Beginning at the Beginning. *Topoi*, 41(5), 939–953. <https://doi.org/10.1007/s11245-022-09797-y>
- Macklin, R. (1981). Can Future Generations Correctly Be Said to Have Rights? In E. Partridge (Ed.), *Responsibilities to future generations: Environmental ethics*. Prometheus Books.
- Maynes, J., & Gross, S. (2013). Linguistic Intuitions. *Philosophy Compass*, 8(8), 714–730. <https://doi.org/10.1111/phc3.12052>
- McAllister, B. (2023). *Seemings and the Foundations of Justification: A Defense of Phenomenal Conservatism* (1st ed.). Routledge.  
<https://doi.org/10.4324/9781003281108>
- Milanović, M., & Milanović, A. (2010). Softening the Words of Crisis: Hedging in Economics- Related Articles Published on the Internet. *Ekonomski Horizonti*, 12, 121–129.
- Millson, J. A. (2021). Accepting & Rejecting Questions: First Steps Toward a Bilateralism for Erotetic Logic. In M. Cordes (Ed.), *Asking and answering: Rivalling approaches to interrogative methods*. Narr Francke Attempto.
- Oliva, T. A., Oliver, R. L., & MacMillan, I. C. (1992). A Catastrophe Model for Developing Service Satisfaction Strategies. *Journal of Marketing*, 56(3), 83–95.  
<https://doi.org/10.1177/002224299205600306>
- Parfit, D. (2011). *On What Matters: Volume I*. Oxford University Press.
- Peirce, C. S. (1902). Vague. In J. M. Baldwin (Ed.), *Dictionary of Philosophy and Psychology*. Macmillan.
- Pravato, G. (2022). Indeterminacy and Normativity. *Erkenntnis*, 87(5), 2119–2141.  
<https://doi.org/10.1007/s10670-020-00293-6>
- Priest, G. (1987). *In Contradiction: A Study of the Transconsistent* (Vol. 25, Issue 3, pp. 380–383). Oxford University Press.

- Pummer, T. (2022). Sorites on What Matters. In J. McMahan, T. Campbell, J. Goodrich, & K. Ramakrishnan (Eds.), *Ethics and Existence* (pp. 498–524). Oxford University Press. <https://doi.org/10.1093/oso/9780192894250.003.0019>
- Rabinowicz, W. (2009). Incommensurability and Vagueness. *Aristotelian Society Supplementary Volume*, 83(1), 71–94. <https://doi.org/10.1111/j.1467-8349.2009.00173.x>
- Raffman, D. (2013). *Unruly Words: A Study of Vague Language*. Oup Usa.
- Rawls, J. (1971). *A Theory of Justice*. Belknap Press of Harvard University Press.
- Raz, J. (2009). *The Morality of Freedom* (Repr). Clarendon Press.
- Razinsky, H. (2016). *Ambivalence: A Philosophical Exploration*. Rowman & Littlefield International.
- Rheinberger, H.-J. (2011). Infra-Experimentality: From Traces to Data, from Data to Patterning Facts. *History of Science*, 49(3), 337–348. <https://doi.org/10.1177/007327531104900306>
- Sainsbury, R. M. (1996). Concepts Without Boundaries. In R. Keefe & P. Smith (Eds.), *Vagueness: A reader*. MIT Press.
- Schauer, F., & Zeckhauser, R. J. (2007). Paltering. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.832634>
- Schiffer, S. (2000). Vagueness and Partial Belief. *Nous*, 34(s1), 220–257. <https://doi.org/10.1111/0029-4624.34.s1.22>
- Schiffer, S. (2003). *The Things We Mean* (1st ed.). Oxford University Press Oxford. <https://doi.org/10.1093/0199257760.001.0001>
- Schiffer, S. (2020). Quandary and Intuitionism: Crispin Wright on Vagueness. In A. Miller (Ed.), *Logic, Language, and Mathematics* (1st ed., pp. 154–174). Oxford University Press. <https://doi.org/10.1093/oso/9780199278343.003.0007>
- Schoenfield, M. (2016). Moral Vagueness is Ontic Vagueness. *Ethics*, 126(2), 257–282. <https://doi.org/10.1086/683541>
- Searle, J. R. (1975). A Taxonomy of Illocutionary Acts. In K. Gunderson (Ed.), *Language, Mind and Knowledge* (pp. 344–369). University of Minnesota Press.
- Shafer-Landau, R. (1994). Ethical Disagreement, Ethical Objectivism and Moral Indeterminacy. *Philosophy and Phenomenological Research*, 54(2), 331–344.
- Shafer-Landau, R. (1995). Vagueness, Borderline Cases and Moral Realism. *American Philosophical Quarterly*, Vol. 32, No. 1, 83–96.
- Shapiro, S. (2006). *Vagueness in Context* (Vol. 76, Issue 2, pp. 471–483). Oxford

- University Press.
- Smith, N. J. (2005). Vagueness as Closeness. *Australasian Journal of Philosophy*, 83(2), 157–183. <https://doi.org/10.1080/00048400500110826>
- Sorensen, R. (2011). Vague Music. *Philosophy*, 86(2), 231–248. <https://doi.org/10.1017/S0031819111000052>
- Sorensen, R. (2023). Vagueness. *The Stanford Encyclopedia of Philosophy (Winter 2023 Edition)*. <<https://plato.stanford.edu/archives/sum2022/entries/vagueness/>>.
- Sorensen, R. A. (1988). *Blindspots*. Clarendon Press ; Oxford University Press.
- Sorensen, R. A. (1991). Moral Dilemmas, Thought Experiments, and Conflict Vagueness. *Philosophical Studies*, 63(3), 291–308. <https://doi.org/10.1007/BF00354195>
- Sorensen, R. A. (2008). *Vagueness and Contradiction* (Reprint). Clarendon Press.
- Sosa, E. (2007). *A Virtue Epistemology* (Vol. 143, Issue 3, pp. 427–440). Oxford University Press.
- Steup, M. (2023). The Chemistry of Epistemic Justification. In K. McCain, S. Stapleford, & M. Steup, *Seemings* (1st ed., pp. 7–22). Routledge. <https://doi.org/10.4324/9781003299349-3>
- Stöttinger, E., Sepahvand, N. M., Danckert, J., & Anderson, B. (2016). Assessing Perceptual Change With an Ambiguous Figures Task: Normative Data for 40 Standard Picture Sets. *Behavior Research Methods*, 48(1), 201–222. <https://doi.org/10.3758/s13428-015-0564-5>
- Sud, R. (2019). Moral Vagueness as Semantic Vagueness. *Ethics*, 129(4), 684–705. <https://doi.org/10.1086/702976>
- Tappolet, C. (2013). Evaluative vs. Deontic Concepts. In H. LaFollette (Ed.), *International Encyclopedia of Ethics* (p. wbiee118). Blackwell Publishing Ltd. <https://doi.org/10.1002/9781444367072.wbiee118>
- Taylor, D. E. (2018). *A Minimal Characterization of Indeterminacy*. 18(5).
- Tesser, A., & Achee, J. (1994). Aggression, Love, Conformity, and Other Social Psychological Catastrophes. In R. R. Vallacher & A. Nowak (Eds.), *Dynamical systems in social psychology*. Academic Press.
- Townsend, D., & Bever, T. G. (2001). *Sentence Comprehension: The Integration of Habits and Rules*. The MIT Press. <https://doi.org/10.7551/mitpress/6184.001.0001>
- Tucker, C. (2013). Seemings and Justification: An Introduction. In C. Tucker (Ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal*

- Conservatism* (pp. 1–29). Oxford University Press.
- Tye, M. (1994). Sorites Paradoxes and the Semantics of Vagueness. *Philosophical Perspectives*, 8, 189. <https://doi.org/10.2307/2214170>
- Van Der Maas, H. L. J., Kolstein, R., & Van Der Pligt, J. (2003). Sudden Transitions in Attitudes. *Sociological Methods & Research*, 32(2), 125–152. <https://doi.org/10.1177/0049124103253773>
- Weatherson, B. (2004). *Three Objections to Smith on Vagueness*. (manuscript). <https://philpapers.org/archive/WEATOT.pdf>
- Weatherson, B. (2010). Vagueness as Indeterminacy. In R. Dietz & S. Moruzzi (Eds.), *Cuts and Clouds: Vagueness, its Nature and its Logic*. Oxford University Press.
- Williams, J. R. G. (2012). Indeterminacy and Normative Silence. *Analysis*, 72(2), 217–225.
- Williams, J. R. G. (2016). Indeterminacy, Angst and Conflicting Values. *Ratio*, 29(4), 412–433. <https://doi.org/10.1111/rati.12141>
- Williams, J. R. G. (2017). Indeterminate Oughts. *Ethics*, 127(3), 645–673.
- Williamson, T. (1996). *Vagueness*. Routledge.
- Williamson, T. (2007). *The Philosophy of Philosophy* (Issue 3). Wiley-Blackwell.
- Woodward, J. (2010). Data, Phenomena, Signal, and Noise. *Philosophy of Science*, 77(5), 792–803. <https://doi.org/10.1086/656554>
- Wright, C. (1975). On the Coherence of Vague Predicates. *Synthese*, 30(3–4), 325–65.
- Wright, C. (2001). On Being in a Quandary. Relativism Vagueness Logical Revisionism. *Mind*, 110(437), 45–98. <https://doi.org/10.1093/mind/110.437.45>
- Zeller, K. (2024). Seemings: Justifying Beliefs Based on How Things Seem. *1000 Word Philosophy*. [1000wordphilosophy.com/2024/01/16/seemings](https://1000wordphilosophy.com/2024/01/16/seemings)