

# **Vagueness in What Matters**

Superhard Comparisons, Choice and Vagueness

Miguel F. Dos Santos  
Department of Philosophy  
UCL

A thesis submitted for the degree of  
Doctor of Philosophy

2019

## Declaration

I, Miguel Fernandes Alves Ferreira Dos Santos, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

---

## Abstract

What is the minimum number of grains that you need to remove from a heap of sand for it to no longer be a heap? Is a male adult who is 180 cm in height tall? Over the past few decades, the philosophical literature has appealed to cases such as these, which exemplify the vagueness of ‘is a heap’ and ‘is tall’ (respectively), to solve problems in the metaphysics, logic and epistemology of vagueness. The driving thought of this dissertation is that vagueness can arise in the application of much more interesting predicates than ‘is a heap’ and ‘is tall’. In particular, vagueness can arise in the application of predicates that matter to us, such as ‘is better than’ and ‘is morally permissible’. Unsurprising as this thought may seem, it allows me to explore uncharted territory. In Chapter 1, I develop an account of decision-making under vagueness in ‘better than’ with the aim of coming up with a novel solution to the problem of superhard comparisons, the problem of explaining what grounds cases of comparison in which (roughly) no classic value relation—‘better than’, ‘worse than’ or ‘equally good’—appears to hold between a certain pair of items relative to a value that they bear. In Chapter 2, I link my approach to the problem of superhard comparisons with the growing literature on moral vagueness, arguing that different leading theories of vagueness yield different accounts of decision-making under moral vagueness. In Chapter 3, building on the approach and conclusions of the previous chapters, I develop a solution to a neglected problem in aesthetics.

## Impact statement

This dissertation brings together research areas in philosophy that either have only occasionally been brought together (Chapters 1–2) or have never been brought together at all (Chapter 3). Despite the pervasiveness of vagueness, not much has been said in philosophy about how vagueness interacts with value theory (e.g., with value relations like ‘better than’) and normative theory (e.g., with moral permissibility). More generally, very little has been said about how vagueness might constrain decision-making—the focal point of the dissertation. Taking a step in bringing these areas closer together, this dissertation highlights some of the implications of vagueness for value and normative theories, and, conversely, it highlights some of the implications of evaluative and normative vagueness for theories of vagueness. In addition to being of interest to philosophers working on vagueness (or indeterminacy) as well as to philosophers working in value and normative theories, Chapter 1 should be of special interest to decision theorists and Chapter 3 to art historians. For those outside academia, the main contribution of this dissertation consists in exemplifying the pervasiveness of vagueness in our daily life, especially as a rich but overlooked source of uncertainty about what to do—a source that has the potential to shed new light on some long-standing debates in the public sphere.

## Acknowledgments

The research that led to this dissertation was funded by a Ph.D. scholarship to UCL from the Portuguese Foundation for Science and Technology and the European Union, and by a Fulbright scholarship to Princeton University from the Fulbright Commission. I would like to begin by thanking these institutions for their invaluable financial support. Thank you also to the three Departments of Philosophy in which this research was carried out: UCL, Princeton and Uppsala.

For very helpful feedback on the material of this dissertation, or for valuable discussions about some of the (many) issues with which it deals, I would like to thank: Ulrike Heuer, Sam Liao, Sherif Girgis, Ruth Chang, Robert Williams, Peter Singer, Olle Risberg, Matti Eklund, Lavinia Picollo, Laura Silva, James Wilson, John Hyman, John Broome, Jens Johansson, Folke Tersman, Elisabeth Schellekens, Doug Lavin, Catharine Abell, Christy Mag Uidhir, Brian Ball, Aaron Meskin, and two anonymous referees for the *Australasian of Journal Philosophy* (in which a version of Chapter 1 is forthcoming). Most especially, though, I would like to thank Joe Horton, Gideon Rosen and Erik Carlson.

I have no words to thank my parents, my grandparents and Triin.

Stockholm, November 2019

To Madalena Alves Santos and José Manuel Santos, my parents

# Contents

## INTRODUCTION // 11

## CHAPTER 1: SUPERHARD CHOICES // 15

1. INTRODUCTION // 15
2. ON THE STRUCTURE OF SUPERHARD COMPARISONS // 19
  - 2.1. *The Collapsing argument against Trichotomy-denying views* // 25
  - 2.2. *The vagueness view, according to Broome* // 26
  - 2.3. *The standard objection to Broome's vagueness view* // 27
3. SUPERHARD COMPARISONS AND CHOICE // 28
  - 3.1. *Answering the Top and Middle Area Questions* // 31
  - 3.2. *Decision-making under (asymmetric) vague betterness* // 32
  - 3.3. *The Make No Mistake argument against Trichotomy-denying views* // 34
  - 3.4. *Is Make No Mistake true?* // 34
  - 3.5. *Does Must Not Believe It really follow from supervaluationism?* // 36
  - 3.6. *Could Trichotomy deniers save their views by appealing to higher-order vagueness?* // 37
4. A NEW (INTERPRETATION OF THE) VAGUENESS VIEW // 38
5. THE MAKE NO MISTAKE ARGUMENT, GENERALISED // 44
  - 5.1. *A recipe for generalisation* // 44
  - 5.2. *A theory of vagueness to which the Make No Mistake argument cannot be generalised* // 45
6. DECISION-MAKING UNDER VAGUENESS // 48
  - 6.1. *Decision-making under normative vagueness* // 49
    - 6.1.1. Williams' account // 49
    - 6.1.2. Bales' account // 52
  - 6.2. *Decision-making under doxastic vagueness* // 54
    - 6.2.1. Rinard's account // 54
  - 6.3. *Decision-making under evaluative vagueness* // 57
7. CONCLUSION // 59

## CHAPTER 2: SUPERHARD CHOICES AND MORAL VAGUENESS // 60

1. INTRODUCTION // 60
2. TWO CASES OF MORAL VAGUENESS FROM PRACTICAL ETHICS // 63
  - 2.1. *A case of moral vagueness from the abortion debate: Abortions* // 63

2.1.1.	The simple liberal view //	64
2.1.2.	A moral sorites paradox on the basis of Abortions //	66
2.2.	<i>A case of moral vagueness from the animal welfare debate: Meals</i> //	68
2.2.1.	Simple moral vegetarianism //	69
2.2.2.	A moral sorites paradox on the basis of Meals //	70
2.3.	<i>A common source</i> //	72
2.4.	<i>A common structure</i> //	73
3.	MORAL VAGUENESS AND SYMMETRIC VAGUE BETTERTNESS //	74
3.1.	<i>Francesca's moral vagueness case and symmetric vague betterness</i> //	76
3.1.1.	Late Stage //	76
3.1.2.	Early Stage //	77
3.1.3.	Francesca's is a case of symmetric vague betterness //	77
3.2.	<i>Tomaso's moral vagueness case and symmetric vague betterness</i> //	78
3.2.1.	Chain's End //	78
3.2.2.	Chain's Beginning //	79
3.2.3.	Tomaso's is a case of symmetric vague betterness //	79
3.3.	<i>The Choice Under Symmetric Vague Betterness Question</i> //	80
4.	DECISION-MAKING UNDER SYMMETRIC VAGUE BETTERTNESS //	80
4.1.	<i>Mistakes, moral mistakes, prudential mistakes</i> //	81
4.1.1.	What Francesca also knows about her choice situation //	84
4.1.2.	What Tomaso also knows about his choice situation //	88
4.1.3.	The upshot so far //	91
4.2.	<i>Vagueness-as-ignorance</i> //	92
4.3.	<i>The Play-it-Safe account of decision-making under symmetric vague betterness</i> //	95
4.4.	<i>Vagueness-as-absence-of-truth</i> //	102
4.4.1.	What is indeterminacy? //	103
4.4.2.	Vagueness-as-absence-of-truth, given the minimal theory //	105
4.5.	<i>The Up-to-You account of decision-making under symmetric vague betterness</i> //	108
4.5.1.	The Action Under Normative Indeterminacy Question //	112
5.	IMPLICATIONS FOR DEBATES IN PRACTICAL ETHICS //	115
6.	THE NEW VAGUENESS VIEW, REVISITED //	117
6.1.	The TVG and IG vagueness views //	118
6.2.	The PL vagueness view //	120
7.	CONCLUSION //	124



## **CHAPTER 3: SUPERHARD CHOICES AND INDETERMINATE ART // 125**

### **1. INTRODUCTION // 125**

### **2. THE DUCHAMP'S BOX PROBLEM // 130**

#### *2.1. Duchamp's paradox // 131*

#### *2.2. The theoretical problem and its meta-problem // 135*

#### *2.3. Two problems from art history // 138*

##### *2.3.1. The problem of the change in the practice // 138*

##### *2.3.2. The problem of the two avant-gardes via Dada // 143*

##### *2.3.3. We need a unified solution to these two problems // 152*

#### *2.4. The practical problem and its meta-problem // 152*

#### *2.5. Why do we need a solution to the Duchamp's box problem? // 158*

### **3. THE NORMATIVE INDETERMINACY HYPOTHESIS // 160**

#### *3.1. Duchamp's paradox, revisited // 160*

#### *3.2. Does Duchamp's paradox have a happy-face solution? // 162*

### **4. SHARPENING THE NORMATIVE INDETERMINACY HYPOTHESIS // 165**

#### *4.1. Ásta's conferralist account of social properties // 166*

#### *4.2. The Mutt affair and the origins of the change in the practice, revisited // 171*

#### *4.3. The F-is-art sharpening and the F-is-not-art sharpening // 174*

### **5. A SOLUTION TO THE DUCHAMP'S BOX PROBLEM // 177**

#### *5.1. A unified solution to the First Pair of Problems // 179*

#### *5.2. Action under (unsuspected) normative indeterminacy // 180*

##### *5.2.1. Five features of action under unsuspected normative indeterminacy // 181*

##### *5.2.2. The Mutt affair, revisited—yet again // 184*

#### *5.3. A unified solution to the Second Pair of Problems // 187*

#### *5.4. A unified solution to the Third Pair of Problems // 189*

#### *5.5. A unified solution to all problems // 191*

### **6. AN ALTERNATIVE INTERPRETATION OF THE HYPOTHESIS // 192**

### **7. CONCLUSION // 193**

## **CONCLUSION // 194**

## **BIBLIOGRAPHY // 197**



## Introduction

There is a heap of sand in front of you. What is the minimum number of grains that you need to remove from it for it to no longer be a heap? Brad Pitt is 180 cm in height. Is he tall (for a male adult)? Over the past few decades, the literature on vagueness has appealed to cases such as these, which exemplify the vagueness of ‘is a heap’ and ‘is tall’ (respectively), to answer at least two central questions. One of them is a question in the *metaphysics of vagueness*: what is the nature or source of vagueness? The other is a question in the *logic of vagueness*: what is the logic of vagueness—for instance, is it classical or non-classical?

The driving thought of *Vagueness in What Matters* is that vagueness can arise in the application of much more interesting predicates than ‘is a heap’ and ‘is tall’. In particular,

vagueness can arise in the application of predicates that *matter* to us, such as ‘is better than’ and ‘is morally permissible’. Unsurprising as this thought may seem, it allows me to explore uncharted territory, by asking the following questions:

*Choice Under Vague Betterness Question:*

In a choice situation between a pair of options,  $x$  and  $y$ , what is it that one rationally ought to do if one knows that it is vague whether  $x$  is better than  $y$ ?

*Choice Under Moral Vagueness Question:*

What is it that one rationally ought to do if one knows that it is (determinately) prudentially better to  $\varphi$ , but that it is vague whether it is morally permissible to  $\varphi$ ?

In tackling these two questions, I explore the implications of leading theories of vagueness—most notably, supervaluationism and epistemicism—for rational belief. More specifically, in developing accounts of decision-making under vagueness in ‘better than’ (in response to the former question) and under moral vagueness (in response to the latter question), I consider what was until recently a neglected question in the vagueness literature: which doxastic attitude ought one to have towards  $p$  if one knows that it is vague whether  $p$ ? This is a question in the *epistemology of vagueness*.

The dissertation is divided into three chapters. Chapter 1 (“Superhard choices”) focuses on the problem that is at the core of the dissertation: the problem of so-called *superhard* comparisons. This is the problem of explaining what grounds cases of comparison in which (roughly) no classic value relation—‘better than’, ‘worse than’ or ‘equally good’—appears to hold between a certain pair of items relative to a value that they bear. On the dominant views, held most notably by Joseph Raz [1986] and Ruth Chang [2002a, 2002b], superhard comparisons are grounded, at least partially, in the failure of ‘better than’, ‘worse than’ and ‘equally good’. On an alternative view, which might be called the vagueness view, first developed by John Broome [1997], they are grounded in vagueness about which of the classic

value relations holds between the items. In this chapter of the dissertation, I focus especially on superhard comparisons in the context of choice—superhard choices—and develop a novel argument against the dominant views based on a partial account of decision-making under vagueness in ‘better than’ (or vague betterness, as I prefer to put it). The upshot is that a new interpretation of the vagueness view emerges. I call it the *new vagueness view* of superhard comparisons and choices. Throughout the chapter, supervaluationism is assumed, but later in it I show how to generalise my argument to other theories of vagueness, such as Williamson’s epistemicism.

At the core of the new vagueness view lies a distinction between two kinds of vague betterness. Cases of what I call *symmetric* vague betterness are cases in which (for a pair of items,  $x$  and  $y$ ) it is vague whether  $x$  is better than  $y$ , and vague whether  $y$  is better than  $x$ . On the other hand, cases of what I call *asymmetric* vague betterness are cases in which it is vague whether  $x$  is better than  $y$ , but  $y$  is (determinately) not better than  $x$ . The account of decision-making on the basis of which I argue against the dominant views of superhard comparisons and choices is a *partial* account of decision-making under vague betterness because it is only an account of decision-making under *asymmetric* vague betterness. Accordingly, Chapter 1 does not say much about decision-making under symmetric vague betterness.

Chapter 2 (‘Superhard choices and moral vagueness’) takes up the task of saying something substantial about decision-making under symmetric vague betterness, but for a specific set of such cases, namely, cases of symmetric vague betterness that have their source in vagueness in ‘is morally permissible’. What I do in this chapter of the dissertation is thus to link my approach to the problem of superhard comparisons with the growing literature on moral vagueness, the aim being to develop a more comprehensive answer to the Choice Under Vague Betterness Question and to answer (at least for the cases I consider) the Choice Under Moral Vagueness Question, which, I show, turn out to be closely related enterprises. More specifically, I argue that what I call (following David Barnett [2011]) vagueness-as-ignorance

and vagueness-as-absence-of-truth—of which supervaluationism and some ontic theories of vagueness are versions—yield rival accounts of decision-making under the relevant cases of symmetric vague betterness, a result that carries interesting implications, I argue, for at least two debates in practical ethics: the abortion debate and animal welfare debate.

Given some theories of vagueness, such as vagueness-as-absence-of-truth, vagueness is a kind of indeterminacy, where the latter is a more general phenomenon than vagueness and can be understood, minimally, as follows: if it is indeterminate whether  $p$ , there is no fact of the matter about whether  $p$ . One conclusion of Chapter 2 (conditional on vagueness-as-absence-of-truth) is that sometimes one acts, without knowing or believing that one is acting, in choice situations such that it is indeterminate whether one ought to  $\varphi$  and indeterminate whether one ought not to  $\varphi$ —choice situations in which there is *normative indeterminacy*. Chapter 3 (“Superhard choices and indeterminate art”) applies and extends some of the conclusions of Chapter 2 with the aim of developing a solution to a problem in seemingly unrelated territory: what I call the *Duchamp’s box problem*, the neglected problem of finding a unified solution to a set of hard problems that have Marcel Duchamp’s *Fountain* at their core. To solve the Duchamp’s box problem, I hypothesise that *Fountain* gives rise to normative indeterminacy.

## *Chapter 1*

### Superhard choices

#### 1. INTRODUCTION

How do Shakespeare’s artistic achievements compare to Michelangelo’s? Is Paris more impressive than New York City? In recent years, interest has grown in comparisons like these, cases in which, roughly, no classic value relation—‘better than’, ‘worse than’ or ‘equally good’—appears to hold between a certain pair of items (e.g., Paris and New York City) relative to a value that they bear (e.g., impressiveness). They have been referred to as *superhard comparisons* [Chang 2002a]. What makes them *superhard*? That they falsify at least one of the following claims:

*Trichotomy*: For any pair of items,  $x$  and  $y$ , bearing a value  $V$ , either  $x$  is better than  $y$  relative to  $V$ , or  $x$  is worse than  $y$  relative to  $V$ , or  $x$  and  $y$  are equally good relative to  $V$ .

*Knowability*: For any pair of items,  $x$  and  $y$ , bearing a value  $V$ , it is knowable which of the three classic value relations holds between  $x$  and  $y$  relative to  $V$ .

But helpful as this characterisation may be, it is compatible with a range of views about what grounds superhard comparisons.<sup>1</sup>

According to one of the two dominant accounts of superhard comparisons, held most notably by Joseph Raz [1986], a comparison is superhard because the items being compared are *incomparable* relative to  $V$ : neither is better than the other, nor are they equally good, *all things considered* (relative to  $V$ ). Consider the comparison between the artistic achievements of Shakespeare and those of Michelangelo. Michelangelo's achievements are greater than Shakespeare's in some respects, the former's being, in the words of his contemporary Giorgio Vasari [1550: 282], 'supreme not merely in one of [the] arts but in all three at once', namely, sculpture, painting and architecture. But Shakespeare's achievements are greater in other respects, given his timeless contributions to poetry and theatre, and the fact these were achieved over a fifty-two-year lifespan (against the eighty-eight years of Michelangelo's life).<sup>2</sup> Thus, neither can truly be said to be a better artist than the other, all things considered. By Trichotomy, it follows that they are equally good artists. But it has been argued that they cannot be.

Suppose, for *reductio*, that Shakespeare and Michelangelo are equally good artists. Now suppose that we improve Shakespeare's artistic achievements; for example, by adding to his body of work a play whose significance exceeds that of any of his actual plays. Call this possible artist 'Shakespeare+'. Shakespeare+ is a better artist than Shakespeare—after all, it is, by

---

<sup>1</sup> Why call these comparisons 'superhard' rather than simply 'hard'? Because doing so better individuates the phenomenon I am interested in. To characterise a comparison between  $x$  and  $y$  relative to  $V$  as 'hard' suggests (in most contexts anyway) that it is hard but *not* impossible to come to know which is better than the other or whether they are equally good (relative to  $V$ ). But there appear to be comparisons that, by challenging Trichotomy or Knowability, do not fit this pattern—hence why I prefer to call these superhard comparisons.

<sup>2</sup> Michelangelo was also an accomplished poet, but I ignore this fact here for simplicity's sake.



stipulation, an improved version of him. But if Shakespeare and Michelangelo are equally good artists, then Shakespeare+ must also be better than Michelangelo. Intuitively, however, he is not; slightly improving Shakespeare (resulting in Shakespeare+) is not enough to make him better than Michelangelo. Thus, they are not equally good artists. This ‘small improvement’ argument, as it is known, is the main argument for the *incomparability view* of superhard comparisons.<sup>3</sup>

According to the other dominant account of superhard comparisons, put forward by Ruth Chang [2002a, 2002b], all that the small improvement argument shows is that Shakespeare and Michelangelo are not *trichotomously* comparable, that is, they are not comparable in terms of the three classic value relations. To put it another way, all the argument shows is that Trichotomy is false. But to conclude from this that Shakespeare and Michelangelo are incomparable is to assume that the three classic value relations are exhaustive, that no other value relation can hold between Shakespeare and Michelangelo. In other words, it is to assume that:

*Value Comparability Is Trichotomous:* For any  $x$  and  $y$ , if  $x$  and  $y$  are comparable relative to  $V$ , either  $x$  is better than  $y$  relative to  $V$ , or  $x$  is worse than  $y$  relative to  $V$ , or  $x$  and  $y$  are equally good relative to  $V$ .

It turns out, Chang contends, that Value Comparability Is Trichotomous is false, since there is a neglected, fourth value relation. On this account, Shakespeare and Michelangelo are *on a par*: neither is indeed a better artist than other, nor are they equally good artists, all things considered—as shown by the small improvement argument—but they are still comparable, namely, in terms of the parity relation. Value comparability is *tetrachotomous*. Thus, because

---

<sup>3</sup> The first version of the small improvement argument seems to have been originally developed by Ronald De Sousa [1974], and it was later adapted as an argument for the incomparability view by Raz [1986] as well as Walter Sinnott-Armstrong [1988]. For discussion, see Espinoza [2008].

it is committed to Value Comparability Is Trichotomous, the incomparability view is false. This is the *parity view* of superhard comparisons.<sup>4,5</sup>

On the dominant views, then, Trichotomy is false (and, by implication, so is Knowability). On the incomparability view, Trichotomy is false because a superhard comparison is a case in which no classic value relation holds between  $x$  and  $y$  relative to  $V$ —they are incomparable. On the parity view, Trichotomy is false because a superhard comparison is a case in which a *sui generis* value relation beyond the classic three holds between  $x$  and  $y$  relative to  $V$ —they are on a par. Accordingly, call the dominant views *Trichotomy-denying views*. On an alternative view, first developed by John Broome [1997, 2004], Trichotomy is true but Knowability is false. On (what I call) the *vagueness view*, this is so because a superhard comparison is a case in which it is vague which of the three classic value relations holds between  $x$  and  $y$  relative to  $V$ .<sup>6</sup>

In this chapter, I develop a novel argument against Trichotomy-denying views and, based on it, build the basic framework for a new (interpretation of the) vagueness view. In Section 2, I explore the structure of superhard comparisons by making use of a tool known as ‘a standard configuration’ [Broome 1997]. Doing this allows me to distinguish the above views from each other as well as from views that deny the existence of superhard comparisons. Along

---

<sup>4</sup> Chang’s [2002a] main argument for the parity view is known as the chaining argument. Since the difference between the parity and incomparability views does not matter for my purposes (as I make clear below), I will not discuss it here. But see Elson [2014a] for a vagueness-related objection to it.

<sup>5</sup> Derek Parfit [1984] and James Griffin [1986] proposed what might be thought of as early versions of this view, based on the notion of ‘rough equality’, but they are widely regarded (and rightly so, I believe) to be less compelling than Chang’s parity view.

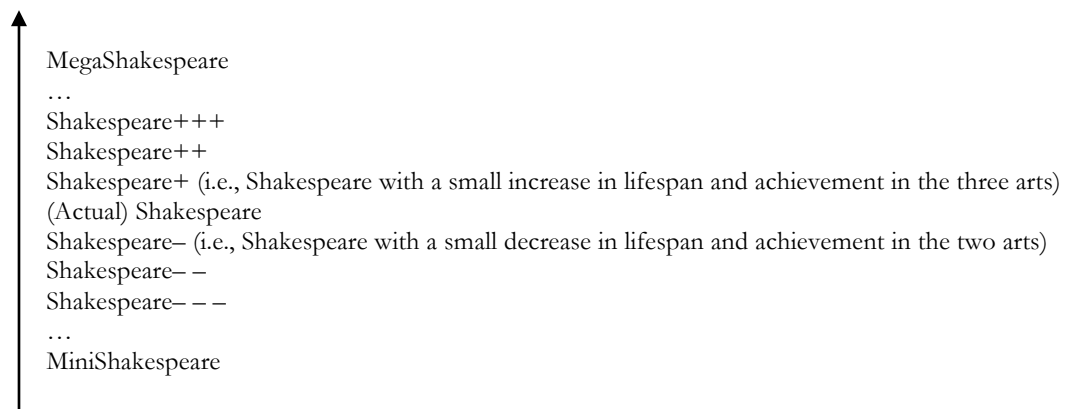
<sup>6</sup> Those familiar with the literature on superhard comparisons may have noticed that I am omitting a view that, even though not popular (but see Regan [1997] for a defence of it), cannot be ruled out without argument—what Chang [2002a] has called ‘epistemicism’. On this view, Shakespeare and Michelangelo are comparable in terms of the classic value relations (relative to the value at stake, artistic goodness), but it is hard, though not impossible, to know which of them holds between them. My reason for not discussing this view here, as an alternative to Trichotomy-denying views and the vagueness view, is twofold. On the one hand, because on my characterisation of the problem of superhard comparisons—in terms of the Trichotomy and Knowability claims—this view turns out to be a *sceptical* view about the existence of superhard comparisons, since it endorses both Trichotomy and Knowability. To put it another way, on this view, superhard comparisons are simply hard comparisons in disguise. (I discuss and reject this view in Section 2.) On the other hand, because I think the most charitable way of understanding this view is as a *version* of (my interpretation of) the vagueness view, namely, a version that is conditional on standard, Williamsonian epistemicism being the correct theory of vagueness. (I discuss this view in Section 5.)

the way, I spell out Broome's vagueness view and the standard objection to it. In Section 3, I turn my attention to superhard comparisons in the context of choice and offer an account of decision-making under vagueness in 'better than', on the basis of which I argue against Trichotomy-denying views. In Section 4, I introduce what I call the *new (interpretation of the) vagueness view* of superhard comparisons. In line with the literature on superhard comparisons, supervaluationism is assumed throughout. In Section 5, however, I provide a recipe for generalising my argument against Trichotomy-denying views beyond supervaluationism, showing that it can be generalised straightforwardly to Williamson's epistemicism. In Section 6, I distinguish my account of decision-making under vagueness from others in the literature.

## 2. ON THE STRUCTURE OF SUPERHARD COMPARISONS


Let us assume that the comparison between Shakespeare and Michelangelo is indeed superhard and consider a variation on Shakespeare's life, a scenario in which his lifespan and artistic achievements are drastically *increased*: he lived as many years as Michelangelo and in those extra years his achievements crossed over into sculpture, painting and architecture to match, or even surpass, Michelangelo's achievements in those arts (the sculpture of David, the frescos on the ceiling of the Sistine Chapel, the Laurentian Library, and so on). Call him 'MegaShakespeare' and the arts of sculpture, painting and architecture (to which MegaShakespeare made astounding contributions, in addition to poetry and theatre) 'the three arts'. In such a scenario, MegaShakespeare is clearly *better than* Michelangelo. Consider now another variation on Shakespeare's life, a scenario in which his lifespan and artistic achievements are drastically *decreased*: he died at (say) the age of twenty and thus never wrote *Sonnet 116*, *Hamlet*, and so on. Call him 'MiniShakespeare' and the arts of poetry and theatre (for which MiniShakespeare had a gift but to which he never made major contributions) 'the two arts'. In such a scenario, MiniShakespeare is clearly *worse than* Michelangelo.

What we have just done in these two variations on Shakespeare's life was to drastically increase or decrease his lifespan and achievements in the two or the three arts, and then compare his (all-things-considered) artistic achievements to Michelangelo's, whose lifespan and achievements in the three arts were held constant. Imagine now a *chain* of variations on Shakespeare's life. At the very top of the chain is MegaShakespeare. Immediately below him is an artist whose lifespan and artistic achievements were a little less drastically increased than MegaShakespeare's; that is, this artist's lifespan and artistic achievements were also drastically increased compared to Shakespeare's (viz., unlike Shakespeare, this artist made astounding contributions to the three arts), but less so than MegaShakespeare's. Immediately below this artist is one whose lifespan and artistic achievements were a little less drastically increased, and so on, all the way down to the very bottom of the chain, where we encounter MiniShakespeare. As we move up along the chain, from MiniShakespeare to MegaShakespeare, every artist is better than every artist that lies below him. Here is a schematic representation of the chain:



Let us call each of the artists on the chain (including MiniShakespeare, MegaShakespeare and the actual Shakespeare) 'a Shakespeare'. Suppose now that we compare each of these Shakespeares to Michelangelo. Michelangelo is not himself part of the chain; rather, he is the standard against which each Shakespeare will be compared. This set of comparisons can

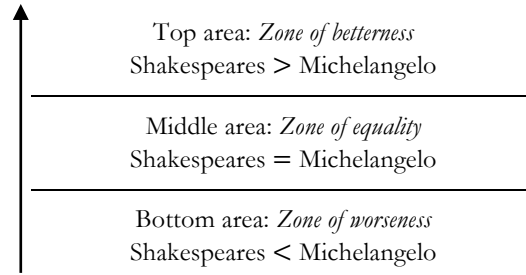
schematically be represented as follows, where ‘>’ stands for ‘better than’, ‘<’ for ‘worse than’, and ‘Shakespeares’ for ‘Shakespeares in this area are’:



S-CONFIGURATION A: Shakespeare vs. Michelangelo case

The Shakespeares in the top area of the chain are better artists than Michelangelo—hence they are in the *zone of betterness*. The Shakespeares in the bottom area of the chain are worse artists than Michelangelo—hence they are in the *zone of worseness*. This chain of Shakespeares, ordered from top to bottom in terms of lifespan and achievements in the two or the three arts, and their comparison to someone who is himself not part of the chain (Michelangelo) is an example of a standard configuration (‘s-configuration’ for short).

Now, what is going on between the zone of betterness and the zone of worseness of S-Configuration A? It depends on which view of superhard comparisons is true. Suppose first that one is sceptical about the existence of superhard comparisons, and so sceptical that the Shakespeare vs. Michelangelo case is a superhard comparison. On (what I call) the *sceptical view*, both Trichotomy and Knowability are true, which entails that either there is nothing going on between the zones of betterness and worseness (there is simply an abrupt transition from one zone to the other) or there is a *zone of equality* between them. Since we are interested in cases, like the Shakespeare vs. Michelangelo case, which suggest that there is a zone of some sort between those two zones, suppose that there is indeed a zone of equality between them (where ‘=’ stands for ‘equally as good as’):



(S-CONFIGURATION) A1: Sceptical view

The zone of equality of A1 has a feature worth highlighting: it has *exactly one* Shakespeare in it. Why? Call ‘Shakespeare-E’ any Shakespeare on the chain who is equally as good as Michelangelo. Take the Shakespeares who are above Shakespeare-E on the chain. Each of these is better than Shakespeare-E (for each Shakespeare on the chain improves on the Shakespeare immediately below). Since Shakespeare-E and Michelangelo are equally good, any Shakespeare above Shakespeare-E must be better than Michelangelo—that is, it must be in the zone of betterness. Take now the Shakespeares who are below Shakespeare-E on the chain. Each of these is worse than Shakespeare-E. Since Shakespeare-E is equally as good as Michelangelo, any Shakespeare below him on the chain must be worse than Michelangelo—that is, it must be in the zone of worseness. Therefore, there is only one Shakespeare (i.e., Shakespeare-E) in the zone of equality.<sup>7</sup>

Is the sceptical view true? There is strong reason to think not, since it cannot accommodate a datum that any view of superhard comparisons should accommodate: ‘better than’ is vague. After all, the sceptical view is committed to there being *no* borderline cases between the Shakespeares who are better (worse) than Michelangelo and the Shakespeare who is equally as good as Michelangelo. In other words, the view is committed to there being a sharp boundary—or a knowable sharp boundary, as epistemicists would have it—between the

<sup>7</sup> I draw here on Broome [1997: 71-2]. As Broome himself points out, this argument assumes the following extension of the principle of transitivity for ‘better than’: ‘For any  $x, y$  and  $z$ , if  $x$  is [better than]  $y$  and  $y$  is equally as [good] as  $z$  or if  $x$  is equally as [good] as  $y$  and  $y$  is [better than]  $z$  then  $x$  is better than  $z$ ’ [Broome 1997: 71].

zone of betterness (worseness) and the zone of equality, as represented by the solid lines dividing the zones in A1. Since it is highly implausible that ‘better than’ is not vague, that it does not admit of borderline cases, the sceptical view is false.

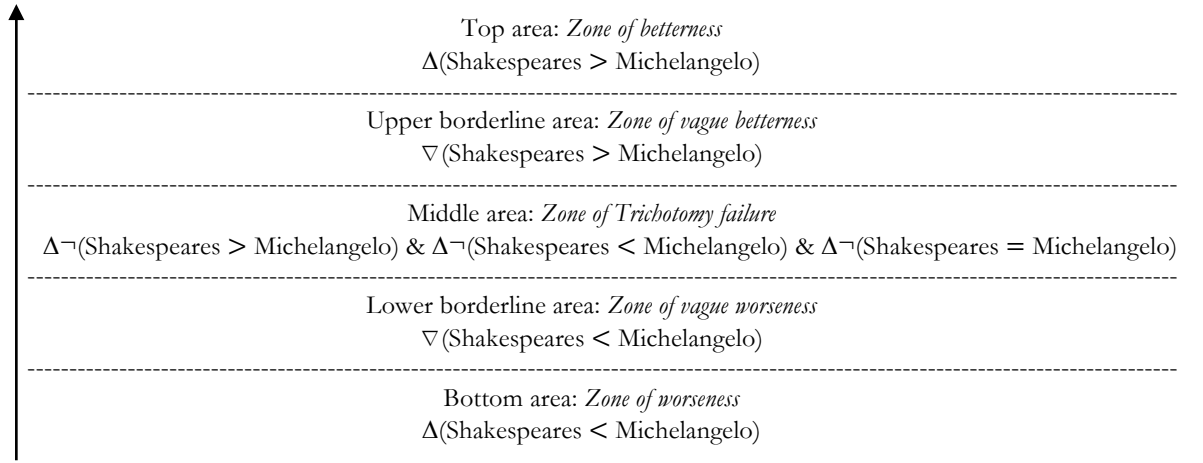
We have just seen, based on a set of variations on the Shakespeare vs. Michelangelo case, that the sceptical view is false because it does not allow for there to be vagueness anywhere between the zone of betterness and the zone of worseness of S-Configuration A. Since the sceptical view is the view that the conjunction of Trichotomy and Knowability is true, it follows that this conjunction is open to counter-examples—in other words, it follows that there are superhard comparisons between the zone of betterness and the zone of worseness of S-Configuration A. With this in mind, we can now ask: what is going on between the zone of betterness and the zone of worseness of S-Configuration A *if* Trichotomy-denying views are true (and, as I have been assuming, the Shakespeare vs. Michelangelo case is a superhard comparison)?

First, there is a *zone of Trichotomy failure*, that is, a zone where no Shakespeare in it is better than, worse than or equally as good as Michelangelo. Is there also a zone of equality, that is, a Shakespeare on the chain that is equally as good as Michelangelo? There cannot be. If there were such a Shakespeare, then, as we saw, any Shakespeare above him on the chain would be in the zone of betterness, and any Shakespeare below him on the chain would be in the zone of worseness. In other words, if there were such a Shakespeare, there would be no zone of Trichotomy failure between the zones of betterness and worseness—there would be a zone of equality instead. So, if Trichotomy-denying views are true, no Shakespeare in S-Configuration A is equally as good as Michelangelo.

Second, and because any view of superhard comparisons must allow for there to be vagueness somewhere between the zones of betterness and worseness of S-Configuration A, there is a *zone of vagueness* between the zone of betterness and the zone of Trichotomy failure: an area, between the top and middle areas of the chain, where it is vague whether the

Shakespeares in it are better than Michelangelo. *Mutatis mutandis*, there is zone of vagueness between the zone of worseness and the zone of Trichotomy failure: an area, between the bottom and middle areas of the chain, where it is vague whether the Shakespeares in it are worse than Michelangelo.

On Trichotomy-denying views, this is then what is going on between the zones of betterness and worseness of S-Configuration A, where ‘ $\nabla$ ’ stands for ‘it is vague whether’, ‘ $\Delta$ ’ for ‘it is determinate that’, and ‘---’ (the dotted lines dividing the zones) for vagueness about where to draw the boundaries of each zone:



(S-CONFIGURATION) A2: Trichotomy-denying views

Are Trichotomy-denying views true? According to Broome [1997], they are false, since A2 is incoherent. To spell out Broome’s argument (and, later, my own argument), I shall from now on assume supervaluationism, according to which truth is truth on all *admissible sharpenings*, where these are acceptable ways of making a vague term precise [Fine 1975, Keefe 2003]. To make this clearer, take the predicate ‘is tall’ (for a male adult) and the following ways of making ‘is tall’ precise:

Sharpening S1: A male adult is tall iff he is 178 cm tall or above.

Sharpening S2: A male adult is tall iff he is 180 cm tall or above.



Sharpening S3: A male adult is tall iff he is 182 cm tall or above.

Each of these sharpenings represents an acceptable way of drawing the boundary between the men who are tall and those who are not tall, and none takes primacy over another. Consider the propositions *Michael Jordan is tall* and *Woody Allen is tall*. Because it comes out true on sharpenings S1, S2 and S3 (i.e., on all admissible sharpenings of ‘is tall’), *Michael Jordan is tall* is true. Because it comes out false on all admissible sharpenings, *Woody Allen is tall* is false. What about the proposition *Brad Pitt is tall*? Is it true or is it false? (He is 180 cm tall.) According to supervaluationism, it is *neither true nor false*, since it is true on some admissible sharpenings (S1 and S2) but false on others (S3).

### 2.1. *The Collapsing argument against Trichotomy-denying views*

Take any Shakespeare in the upper borderline area of A2. Call him ‘Shakespeare-UBORDER’. First, note that, as A2 states, it is neither true nor false that Shakespeare-UBORDER is better than Michelangelo (since this is so for any Shakespeare in the upper borderline area). Second, note that it is false that Michelangelo is better than Shakespeare-UBORDER (i.e., it is false that Shakespeare-UBORDER is worse than Michelangelo). Why? Because this is false for any Shakespeare in the zone of Trichotomy failure and above it. Therefore, these are the structural features of the upper borderline area: it is false that Michelangelo is better than Shakespeare-UBORDER, and neither true nor false that Shakespeare-UBORDER is better than Michelangelo. Now, Broome thinks that this principle holds:

*Collapsing.* For any  $x$  and  $y$ , if it is false that  $y$  is better than  $x$  and not false that  $x$  is better than  $y$ , then it is true that  $x$  is better than  $y$ .

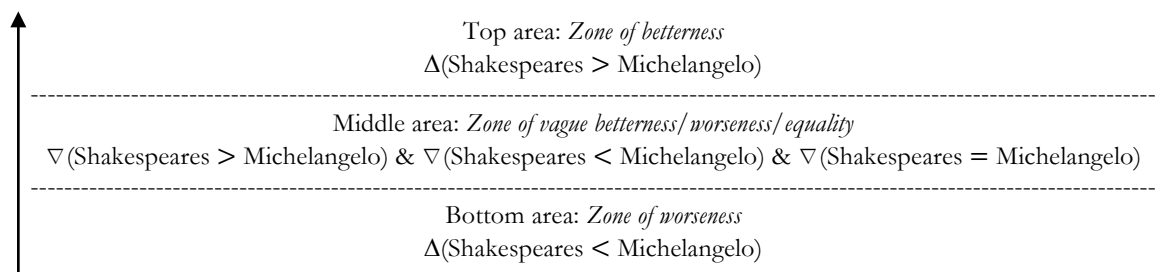
This is his argument for it:

If it is false that  $y$  is [better] than  $x$ , and not false that  $x$  is [better] than  $y$ , then  $x$  has a clear advantage over  $y$  in respect of its [betterness]. So it must be [better] than  $y$ . It takes only the slightest asymmetry to make it the case that one thing is [better] than another. One object is heavier than another if the scales tip ever so slightly toward it. Here there is a clear asymmetry between  $x$  and  $y$  in respect of their [betterness]. That is enough to determine that  $x$  is [better] than  $y$ . [Broome 1997: 74.]

By Collapsing, it follows that, since it is *false* that Michelangelo is better than Shakespeare-UBORDER and *not false* that Shakespeare-UBORDER is better than Michelangelo, it is *true* that Shakespeare-UBORDER is better than Michelangelo. In other words, Shakespeare-UBORDER is in the zone of betterness. This implies that there cannot be a zone of vague betterness after all, since *any* Shakespeare in it ‘collapses’ into the zone of betterness. (*Mutatis mutandis* for the zone of vague worseness.) Thus, to be coherent, A2 cannot have zones of vagueness. But in that case Trichotomy-denying views cannot accommodate the datum that ‘better than’ is vague. Therefore, Trichotomy-denying views are false.

## 2.2. The vagueness view, according to Broome

What is going on between the zones of betterness and worseness of S-Configuration A *if* the vagueness view is true? According to Broome, there is a *single* zone of vagueness, where it is vague which of the three classic value relations holds between the Shakespeares in it and Michelangelo:



(S-CONFIGURATION) A3: Broome's vagueness view

Why are there no upper and lower borderline areas? Because, as we saw, Collapsing rules them out.

### 2.3. *The standard objection to Broome's vagueness view*

Is Broome's vagueness view true? The standard objection to it, developed by Erik Carlson [2004], is that Collapsing is false. To see this, note first that Collapsing is logically equivalent to:

*Symmetry of Vague Betterness*: For any  $x$  and  $y$ , if it is neither true nor false that  $x$  is better than  $y$ , then it is neither true nor false that  $y$  is better than  $x$ .

That Collapsing implies this principle is acknowledged by Broome [1997: 76] himself.<sup>8</sup> But why does the latter imply the former? Notice what would falsify Collapsing: a case in which it is false that  $y$  is better than  $x$ , and neither true nor false that  $x$  is better than  $y$ . But this case would also falsify Symmetry of Vague Betterness, since its being neither true nor false that  $x$  is better than  $y$  is sufficient, by this principle, for its being neither true nor false that  $y$  is better than  $x$ . So, any counter-example to Collapsing is also a counter-example to Symmetry of Vague Betterness. By contraposition, Symmetry of Vague Betterness implies Collapsing.

Now, Symmetry of Vague Betterness is open to a compelling type of counter-example, of which the following case is the best-known instance:

*Alf vs. Beth*: You are comparing two philosophers, Alf and Beth. So far, having considered every property that contributes to being a good philosopher, you have been able to establish that they are equally good, since they possess each of the relevant properties to the same

---

<sup>8</sup> *Proof*. Assume, for *reductio*, that it is true that  $y$  is better than  $x$ , and neither true nor false that  $x$  is better than  $y$ . Then, by the asymmetry of 'better than', it follows that it is *false* that  $x$  is better than  $y$ , which contradicts the assumption. Assume now, also for *reductio*, that it is false that  $y$  is better than  $x$ , and neither true nor false that  $x$  is better than  $y$ . Then, by Collapsing, it follows that it is *true* that  $x$  is better than  $y$ , which contradicts the assumption.

degree. But you have just found out that Alf has greater rhetorical skill than Beth, and wonder whether this makes him a better philosopher than her. Does it? There is no (determinate) answer to this question; it is vague whether possessing rhetorical skill makes him a better philosopher. So, *it is neither true nor false that Alf is a better philosopher than Beth*. However, this much is clear: possessing rhetorical skill does not make him a *worse* philosopher. So, *it is false that Beth is a better philosopher than Alf*.

As should be clear, the conjunction of these two claims—it is false that Beth is a better philosopher than Alf, and neither true nor false that Alf is a better philosopher than Beth—violates Symmetry of Vague Betterness and thereby Collapsing. Broome’s vagueness view is false.<sup>9</sup>

### 3. SUPERHARD COMPARISONS AND CHOICE

Consider the following examples of superhard comparisons:

*Elderly Patient*: You are a physician in a busy hospital and need to decide whether to operate on an elderly patient who has just been admitted into your unit. If you decide to not operate on the patient, she will die, but she will die peacefully. If you decide to operate on her, she will live a few more years but experience intermittent, yet intense, pain. To make the situation more complicated, the patient has no relatives who can help you decide and she is incompetent to decide herself. What should you do?<sup>10</sup>

*Careers*: You are on the verge of making a life-changing decision: to pursue a career in law or a career in music. Your chances of success are the same in both careers, and you are as passionate about law as you are about music. But the music career offers something the law career does

---

<sup>9</sup> Broome [2004: 285-6, n. 6] recognises that the Alf vs. Beth case is a strong objection to the Collapsing argument, but has responded to it by pointing out that one can reasonably doubt the existence of the properties it appeals to—properties (e.g., having rhetorical skill) about which it is vague whether they positively contribute to goodness (e.g., goodness as a philosopher)—and that the objection merely assumes that they do [Broome 2009: 416-7]. However, Broome’s response is unconvincing in several ways, not least because Johan Gustafsson [2013: 436] has shown that the Alf vs. Beth case can be modified so that it does not rely on those properties. Moreover, Luke Elson [2014b] also recently put forward a new counter-example to Collapsing (one that is structurally like Gustafsson’s version of the Alf vs. Beth case). On the other hand, other attempts to save Collapsing from the Alf vs. Beth case (viz., Cristian Constantinescu’s [2012]) have proved unconvincing (see Carlson [2013]). Carlson’s Alf vs. Beth case is here to stay as the standard objection to the Collapsing argument.

<sup>10</sup> Adapted from Sinnott-Armstrong [1988: 66].

not: the possibility of touring the world, with plenty of free time to explore different cultures and places. On the other hand, the law career offers something the music career does not: a high and stable income. What should you do?<sup>11</sup>

*Sartre's Student:* During World War II, a student of Sartre's had a tough decision to make: to stay in France to look after his mother or to join the Free French Forces in England. To leave for England would allow him to promote the values of liberty, equality and fraternity, immortalised in France's national motto. On the other hand, after his father's conviction for treason and his brother's death at the hands of the Nazis, his mother genuinely needed him to stay. What should Sartre's student have done?<sup>12</sup>

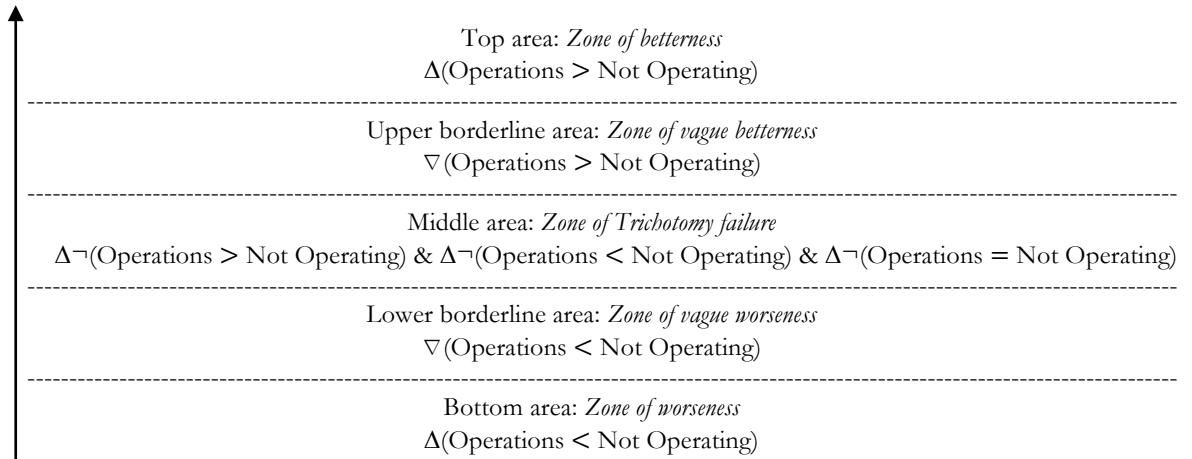
As in the Shakespeare vs. Michelangelo case, we can create an s-configuration for each of these cases. What is distinctive about these cases, in contrast with Shakespeare vs. Michelangelo-type cases, is that we compare with the aim of choosing—they are superhard *choices*. Accordingly, for each area of the s-configuration, we can ask of any option in that area whether it rationally ought to be chosen over the option not on the chain. Let me expand on this point.

Take the Elderly Patient case and assume that the comparison between the option of carrying out the operation (call it 'Operation') and that of not so doing ('Not Operating') is indeed superhard. We can create an s-configuration for it by imagining a chain of Operations—in which the Operation at the very top results in pain that is minute in intensity and frequency, and the one at the very bottom results in excruciating and extremely frequent pain—and by comparing every Operation on the chain to Not Operating. On Trichotomy-denying views, this is the structure of the s-configuration (where 'Operations' stands for 'Operations in this area are'):

---

<sup>11</sup> Adapted from Raz [1986: 126].

<sup>12</sup> Adapted from Broome [2004: 165].



(S-CONFIGURATION) B1: Trichotomy-denying views. Elderly Patient case

(Since there is strong reason to think that Collapsing is false, there is no reason to think that B1 is incoherent.) B1 tells us how any Operation in each area compares to Not Operating. Based on it, we can ask, *for each area*, the following question: what is it that one rationally ought to do if one knows that

(*Top Area Question:*)

Operation-TOP is better than Not Operating?

(*Upper Borderline Area Question:*)

it is vague whether Operation-UBORDER is better than Not Operating?

(*Middle Area Question:*)

Operation-MIDDLE is neither better nor worse than, nor equally as good as, Not Operating?

(*Mutatis mutandis* for the remaining areas.)

The Upper Borderline Area Question asks for something unexplored in the literature (including the vagueness literature): an account of decision-making under vagueness in ‘better than’ or, as I prefer to put it, under vague betterness.<sup>13</sup> In what follows, I provide such an

<sup>13</sup> Robert Williams [2014a, 2016, 2017], Adam Bales [2017] and Susanna Rinard [2015] have developed accounts of decision-making under vagueness, but their accounts are not accounts of decision-making under vagueness of the same type as mine. For more on this, see Section 6.

account, based on which I argue that Trichotomy-denying views are committed to (what I call) a *normative* s-configuration that leaves no room for vagueness, understood supervaluationally, in practical rationality predicates.

### 3.1. *Answering the Top and Middle Area Questions*

To answer our questions, I assume that, when choosing between a pair of items,  $x$  and  $y$ , agents do not make a *mistake in choice* iff  $x$  is not worse than  $y$  and they choose  $x$ . Note that  $x$ 's not being worse than  $y$  is consistent with  $x$ 's being better than  $y$ ,  $x$  and  $y$ 's being equally good,  $x$  and  $y$ 's being incomparable, and  $x$  and  $y$ 's being on a par. The assumption is thus consistent with Trichotomy-denying views, and moreover it is independently plausible [Sen 1997: 746]. In addition to this, I assume that the relevant agents know what is going on in each area of B1 (e.g., they know that, in the top area, Operation-TOP is better than Not Operating).

Take the Top Area Question first. Since agents know that Operation-TOP is better than Not Operating, they ought to believe the proposition *Operation-TOP is not worse than Not Operating*, and ought not to believe the proposition *Not Operating is not worse than Operation-TOP*. This means that—if agents are rational, which I am assuming to be the case throughout—agents believe that they are not making a mistake in choice if they choose Operation-TOP (after all, they know it is *not* worse than Not Operating), and do not believe that they are not making a mistake in choice if they choose Not Operating (they know it *is* worse than Operation-TOP). Thus, Operation-TOP rationally ought to be chosen over Not Operating.<sup>14</sup> One can make sense of this by way of a plausible principle:

*Make No Mistake:* For any options  $x$  and  $y$ , if one ought to believe that  $x$  is not worse than  $y$  and ought not to believe that  $y$  is not worse than  $x$ , then one rationally ought to choose  $x$  over  $y$ .

---

<sup>14</sup> *Mutatis mutandis* for the bottom area—Operation-BOTTOM rationally ought *not* to be chosen.

(Note that this is not to defend Make No Mistake, since one can also make sense of that by way of at least one other plausible principle: for any options  $x$  and  $y$ , if one ought to believe that  $x$  is better than  $y$ , then one rationally ought to choose  $x$  over  $y$ . Rather, this is merely to point out its plausibility. I defend Make No Mistake later, in Section 3.4.)

Consider now the Middle Area Question. Given what agents know—namely, that Operation-MIDDLE is neither better nor worse than, nor equally as good as, Not Operating—it follows that they ought to believe both *Operation-MIDDLE is not worse than Not Operating* and *Not Operating is not worse than Operation-MIDDLE*. In that case, it seems clear that it is rationally permissible to choose either option. In fact, this is what the leading Trichotomy deniers have argued their views to imply for rational choice: when faced with options that are incomparable or on a par, either may rationally be chosen.<sup>15,16</sup> One can make sense of this by way of another plausible principle:

*Cannot Go Wrong Either Way*: For any options  $x$  and  $y$ , if one ought to believe that  $x$  is not worse than  $y$  and ought to believe that  $y$  is not worse than  $x$ , then it is rationally permissible for one to choose  $x$  and rationally permissible for one to choose  $y$ .

### 3.2. *Decision-making under (asymmetric) vague betterness*

What is it that one rationally ought to do if one knows that it is vague whether Operation-UBORDER is better than Not Operating? This is the Upper Borderline Area Question. To answer it, recall that this is also a structural feature of the upper borderline area: Not Operating is (determinately) *not* better than Operation-UBORDER. For this reason, call any case in the

---

<sup>15</sup> See Raz [1997, 1999] and Chang [2005], respectively. Note, however, that Chang [2012, 2016] has argued, contra Raz, that a pair of items' being incomparable precludes rational choice between them.

<sup>16</sup> Given this, one might wonder whether there is any difference in practical decision-making between cases of equality, in which it is also rationally permissible to choose either option, and cases of incomparability and parity. Broome [1999: 154-7] and Chang [2005: 344-7] think that the difference becomes clear when we consider series of choices.



upper borderline area a case of *asymmetric vague betterness*. Because agents know they are before such a case, they know that *Operation-UBORDER is not worse than Not Operating* and that it is vague whether *Not Operating is not worse than Operation-UBORDER*. Clearly, they ought to believe the former proposition. But which doxastic attitude ought they have towards the latter?

Given supervaluationism, *Not Operating is not worse than Operation-UBORDER* is not true. By a truth norm on belief,

*Truth Norm of Belief (TNB)*: One may: believe  $p$  iff  $p$  is true,

it follows that agents ought *not* to believe the proposition. In other words, together with TNB, supervaluationism entails:

*Must Not Believe It*: If one knows that it is vague whether  $p$ , then one ought not to believe  $p$  and one ought not to believe not- $p$ .<sup>17</sup>

Upshot: agents ought to believe *Operation-UBORDER is not worse than Not Operating* and ought *not* to believe *Not Operating is not worse than Operation-UBORDER*. By Make No Mistake, it follows that Operation-UBORDER rationally *ought* to be chosen over Not Operating.<sup>18</sup>

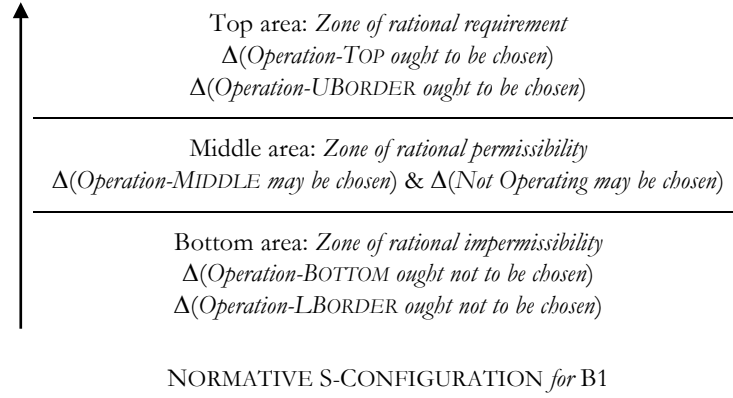
### 3.3. *The Make No Mistake argument against Trichotomy-denying views*

If the argument in Sections 3.1–2 is sound, Trichotomy-denying views are committed to the following *normative s*-configuration:

---

<sup>17</sup> In making this argument, I draw here on recent work by Williams [2014b: esp. 390-6] on the implications of supervaluationism for rational belief. See also Field [2000].

<sup>18</sup> *Mutatis mutandis* for the lower borderline area—Operation-LBORDER rationally ought *not* to be chosen.



Why is this a problem? Because it commits them to the claim that there are no borderline cases between the zone of rational requirement and the zone of rational permissibility, and between the latter and the zone of rational impermissibility. Just as the vagueness of ‘better than’ is a datum that any view of superhard comparisons should accommodate, so too is the vagueness of ‘rationally ought (not) to be chosen over’. Trichotomy-denying views are false.<sup>19</sup>

### 3.4. *Is Make No Mistake true?*

One might object that Make No Mistake is in tension with a plausible alternative principle:

*Not True That One Must Choose It.* For any options  $x$  and  $y$ , if one ought not to believe that  $x$  is better than  $y$ , then it is not true that one rationally ought to choose  $x$  over  $y$ .

If this principle holds, then the right account of decision-making under asymmetric vague betterness is one according to which, to use our case, it is not true that rational agents ought to choose Operation-UBORDER over Not Operating. In other words, there is no reason to doubt that the vagueness in ‘better than’ in the upper borderline area of B1 percolates up, as

<sup>19</sup> It is worth pointing out that for those who wish to deny that ‘rationally ought (not) to be chosen over’ is vague, the burden of proof lies with them, since one can create sorites paradoxes with practical rationality predicates. That these predicates are sorites-susceptible is strong evidence that they are vague. See Horgan et al. [2019].

it were, generating vagueness in ‘rationally ought to be chosen over’ in the corresponding normative s-configuration.

However, Not True That One Must Choose It is open to counter-examples. Consider this case:

*Toys*: Pietro’s daughter does not want just any toy for Christmas: she wants a *red* toy. Pietro promises to give her one. (Assume that breaking the promise is wrong.) However, he leaves it until the last minute to buy the toy and the local Toys ‘R’ Us now offers him only two options: buying a red toy (call this option ‘Red Toy’) or buying a borderline-red toy (‘Borderline-Red Toy’). Pietro is thus in a choice situation in which he knows that Red Toy is *not* worse than Borderline-Red Toy and that it is *vague* whether Borderline-Red Toy is not worse than Red Toy. Which toy ought he to buy?

By Not True That One Must Choose It, it is not true that Pietro rationally ought to choose Red Toy. But this is counter-intuitive. It *is* true that Pietro rationally ought to choose Red Toy, and this is what Make No Mistake entails. Consider now another case, a variation on Toys:

*Toys\**: Everything is the same as in Toys, except that Toys ‘R’ Us offers Pietro these two options instead: buying a black toy (‘Black Toy’) or buying a borderline-red toy (‘Borderline-Red Toy’). Pietro is thus in a choice situation in which he knows that Borderline-Red Toy is *not* worse than Black Toy and that it is *vague* whether Black Toy is not worse than Borderline-Red Toy. Which toy ought he to buy?

Once again, Not True One Must Choose It gives a counter-intuitive verdict, entailing that it is not true that Pietro rationally ought to choose Borderline-Red Toy. Toys and Toys-like cases are thus strong evidence for Make No Mistake.

To be sure, those who lack the intuitions to which I appeal here will no doubt remain unpersuaded by these cases for Make No Mistake, but I expect most people to share the

relevant intuitions—hence why I take Toys and Toys-like cases to be evidence for Make No Mistake.

### 3.5. *Does Must Not Believe It really follow from supervenience?*

One might also object to the argument that supervenience, together with TNB, entails Must Not Believe It. In particular, one might argue that the combination of supervenience and TNB is compatible with another norm about the connection between vagueness and rational belief:

*Vaguely Believe It:* If one knows that it is vague whether  $p$ , then one ought to be such that it is vague whether one believes  $p$ .<sup>20</sup>

If this is right, the argument is flawed.

The response to this objection is that TNB must be distinguished from a closely related norm:

One may: believe  $p$  iff  $p$  is true. (TNB)

One may: believe  $p$  iff  $p$ . (Non-Truth-Centric-NB)

In a classical, non-supervenient setting, TNB and Non-Truth-Centric-NB are equivalent given the T-schema, according to which  $p$  is true iff  $p$ . Under supervenience, however, the claim that  $p$  is true is stronger than the claim that  $p$ —in other words, the T-schema fails. To see why it fails to hold, suppose for reductio that the T-Schema does hold: (1) for any  $p$ ,  $p$  is true iff  $p$ . Given supervenience, vagueness about whether  $p$  entails the following conjunction: (2) it is not true that  $p$ , and it is not true that not- $p$ . By (1) and the first conjunct

---

<sup>20</sup> See Caie [2012] for a defence of a version of Vaguely Believe It. See also Dorr [2003].

of (2), it follows that: (3) not- $p$ . By (1) and the second conjunct of (2), it follows that: (4) not-not- $p$ . By (3) and (4), a contradiction follows: (5) not- $p$  and not-not- $p$ —hence why the T-schema fails to hold under supervaluationism. Accordingly, given supervaluationism, TNB and Non-Truth-Centric-NB are inequivalent and so we must choose which of them to use in theorising.<sup>21</sup>

If we choose TNB, as I did, we get Must Not Believe It as a corollary. If we choose Non-Truth-Centric-NB, we get a norm other than Must Not Believe It as a corollary (e.g., perhaps, Vaguely Believe It). Is the choice between them arbitrary? No, because it is more natural to pair supervaluationism with a *truth* norm on belief, since supervaluationism, at its core, is a claim about truth (viz., that truth is truth on all admissible sharpenings). Accordingly, the choice between those norms is principled and clear: TNB. After all, Non-Truth-Centric-NB does not mention truth at all; it is, rather, a norm that is classically equivalent but supervaluationally inequivalent to a truth norm (viz., TNB). Hence why I chose TNB.

### 3.6. *Could Trichotomy deniers save their views by appealing to higher-order vagueness?*

Assume that ‘better than’ is higher-order vague. Then there is at least one area between the upper borderline area and the middle area of B1 where it is vague whether it is vague whether Operation-UBORDER is better than Not Operating. Call it ‘borderline borderline area’. Since agents do not know, in this area, whether it is vague whether *Not Operating is not worse than Operation-UBORDER*, Must Not Believe It does not kick in to require them not to believe *Not Operating is not worse than Operation-UBORDER*. Thus, Trichotomy deniers might argue that it is *here* that we find vagueness about which option one rationally ought to choose.

To begin with, it is worth pointing out that the assumption that ‘better than’ is higher-order vague is substantive, since it is far from clear that it is a datum that any view of superhard

---

<sup>21</sup> For more on this, precisely in connection to norms, see Williams [2017: esp. 661].

comparisons should accommodate.<sup>22</sup> But even if we grant it, there is reason to doubt that Trichotomy deniers could save their views by appealing to higher-order vagueness. For notice what they would be doing: they would be putting forward an account of decision-making under higher-order vague betterness for cases in which it is vague whether it is vague whether  $x$  is better than  $y$ , *and* false that  $y$  is better than  $x$ . (Why? Because the latter is a structural feature of any case in the middle area of B1 and above it, including the borderline borderline area.) The upshot of it would be that it is vague which of  $x$  and  $y$  rationally ought to be chosen over the other. However, any such account is open to Toys-like counter-examples. Consider this variation on Toys:

*Toys*<sup>#</sup>: Everything is the same as in Toys, except that Toys ‘R’ Us offers Pietro these two options instead: buying a red toy (‘Red Toy’) or buying a borderline-borderline-red toy (‘BB-Red Toy’). Thus, the choice situation is one in which Red Toy is (determinately) *not* worse than BB-Red Toy and it is vague whether it is vague whether BB-Red Toy is not worse than Red Toy. Which toy ought he to buy?

On the Trichotomy deniers’ account of decision-making under consideration, it is vague which toy Pietro ought to buy—and this is implausible.

#### 4. A NEW (INTERPRETATION OF THE) VAGUENESS VIEW

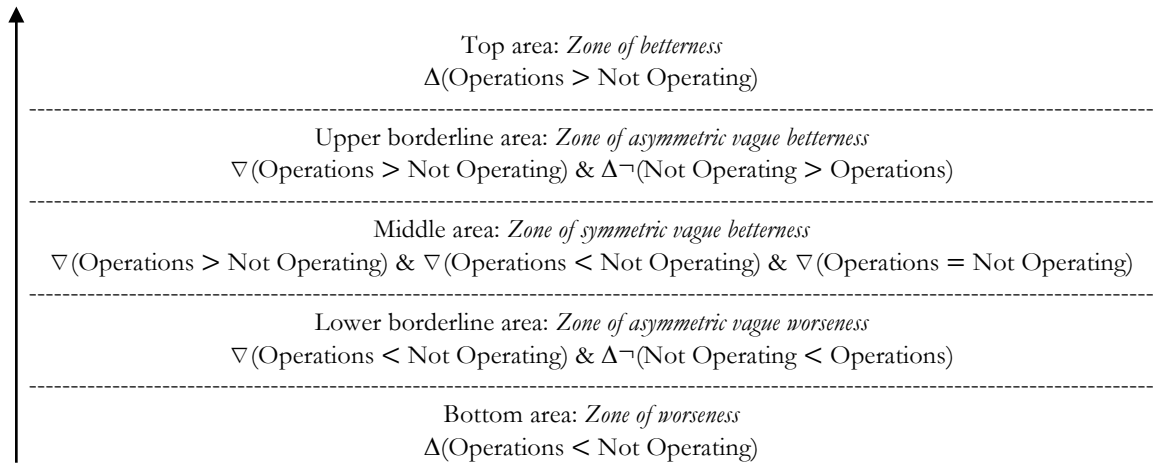
If the Make No Mistake argument is sound, a new interpretation of the vagueness view of superhard comparisons emerges. To see this, recall that, as discussed in Section 2, the Alf vs. Beth case is a counter-example to Symmetry of Vague Betterness. Given this result (due to Carlson), the Alf vs. Beth case is evidence for the following principle:

---

<sup>22</sup> In fact, some proponents of influential theories of vagueness—for instance, Elizabeth Barnes [2010: 618]—argue that it is an advantage of their theories that higher-order vagueness is ruled out.

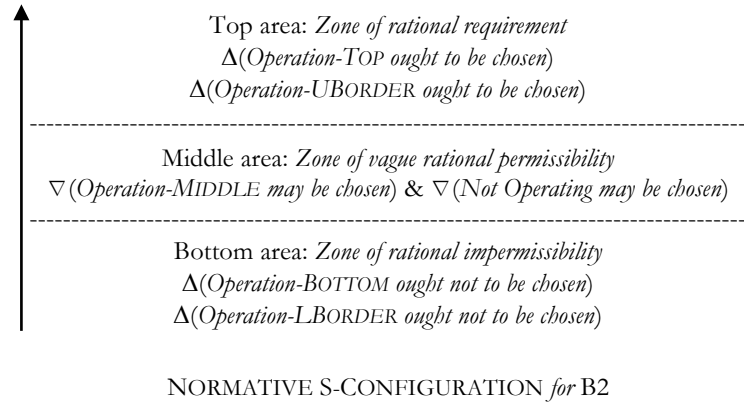
*Non-Symmetry of Vague Betterness*: For any  $x$  and  $y$ , if it is vague whether  $x$  is better than  $y$ , then either  $y$  is (determinately) not better than  $x$  or it is vague whether  $y$  is better than  $x$ .

This is consistent with there being two types of vague betterness cases: those I have called cases of *asymmetric* vague betterness and those I shall call cases of *symmetric* vague betterness, in which it is vague whether  $x$  is better than  $y$  and vague whether  $y$  is better than  $x$ . This yields an interpretation of the vagueness view on which an s-configuration for any superhard comparison is such that it has *three* zones of vagueness (as opposed to one zone, as in Broome's interpretation of the view), as represented below for the Elderly Patient case:



(S-CONFIGURATION) B2: New (interpretation of the) vagueness view. Elderly Patient case

Furthermore, according to this interpretation of the vagueness view, the normative s-configuration corresponding to B2 is such that it has a single zone of vagueness:



Why? First, because, given the account of decision-making under asymmetric vague betterness (worseness) that I defended in Section 3, any case in the upper borderline area of B2 is such that one rationally ought to choose Operation-UBORDER over Not Operating—hence why any such case is in the zone of rational requirement of the normative s-configuration for B2. (*Mutatis mutandis* for any case in the lower borderline area of B2.) Second, because the middle area of B2 raises a question,

*(B2 Middle Area Question)*

What is it that one rationally ought to do if one knows that it is vague whether Operation-MIDDLE is better than Not Operating and vague whether Not Operating is better than Operation-MIDDLE?

to which the answer is, according to the interpretation of the vagueness view that I propose here, that it is vague whether one rationally may choose Operation-MIDDLE, and vague whether one rationally may choose Not Operating. In other words, the interpretation of the vagueness view that I propose here allows for the vagueness in ‘better than’ in the middle area of B2 (i.e., the vagueness in cases of symmetric vague betterness) to percolate up, generating vagueness in rational choice. This is the *new interpretation of the vagueness view*—or simply the *new vagueness view*, as I prefer to call it.



As the latter feature of the new vagueness view suggests, to provide a fully-fledged defence of the view requires that an account of decision-making under *symmetric* vague betterness be developed, namely, an account according to which in *at least some* cases of symmetric vague betterness in the context of choice, like Elderly Patient, there is vagueness about what one rationally may do. This is enough to show that, unlike Trichotomy-denying views, the new vagueness view leaves room for vagueness in practical rationality predicates. To illustrate this feature of the view with a case besides Elderly Patient, consider the following variation on Toys:

*Borderline-Red vs. Borderline-Toy*: Everything is the same as in Toys, except that Toys ‘R’ Us offers Pietro these two options instead: buying a toy that is borderline-red (‘Borderline-Red Toy’) or buying a borderline-toy that is red (‘Red Borderline-Toy’). Thus, the choice situation is one in which it is vague whether Borderline-Red Toy is not worse than Red Borderline-Toy and vague whether Red Borderline-Toy is not worse than Borderline-Red Toy. Which option ought he to choose?

If, as seems plausible, it is vague which option Pietro (who knows the choice situation he is in) rationally ought to choose, then the right account of decision-making under symmetric vague betterness must deem it so. Once such an account is developed, we will have a full account of decision-making under vague betterness (viz., under symmetric and asymmetric vague betterness) and will have provided a fully-fledged defence of the new vagueness view.

With the new vagueness view fleshed out, let me now return to the Alf vs. Beth case. As discussed in Section 2, the Alf vs. Beth case is the standard objection to Broome’s (interpretation of the) vagueness view, but, as I hope is now clear, it is no objection to the new vagueness view. After all, the new vagueness view is consistent with Non-Symmetry of Vague Betterness. However, the Alf vs. Beth case is not merely consistent with the new vagueness

view. It *supports* it, if we assume that Alf and Beth are being compared with the aim of choosing between them. Consider the Alf vs. Beth case in a context of choice:

*Schock Prize*: You are the head of the Royal Swedish Academy of Sciences committee responsible for awarding the 2020 Schock Prize in Logic and Philosophy, and you have narrowed down the list of candidates for the honour to two: Alf and Beth. Having considered every property that contributes to being a good philosopher, you have been able to establish that they are equally good, since they possess each of the relevant properties to the same degree. But you have just found out that Alf has greater rhetorical skill than Beth, and you know that it is vague whether possessing rhetorical skill makes Alf a better philosopher. So, you are now able to establish that *it is vague whether Alf is a better philosopher than Beth*. However, you also know this: possessing rhetorical skill does not make Alf a worse philosopher. So, you can also establish that *Beth is (determinately) not a better philosopher than Alf*. To whom ought you to award the Schock Prize?

I think it is quite clear. Assuming that you know the choice situation you are in, you rationally ought to award the Schock Prize to Alf, since (given what you know) you ought to believe that Alf is not worse than Beth and ought not to believe that Beth is not worse than Alf. The Schock Prize case has the same structure as the cases that vindicate Make No Mistake, namely, any case of asymmetric vague betterness in a context of choice. (As acknowledged before, when Toys and Toys\* were introduced, I do not expect *everyone* to find the verdicts that my view yields about these cases to be intuitive, but I expect most people to find them intuitive.) This makes apparent one of the insights behind the view proposed here: when we shift the Alf vs. Beth case to a context of choice (resulting in Schock Prize), the case turns out to be evidence for (rather than against) the vagueness view, or rather a *new* interpretation of the vagueness view.

Finally, consider Broome's own decision-theoretic case in support of his vagueness view (and more specifically of Collapsing):

*Canberra Suburbs*: ‘Suburbs in Canberra are named after great Australians, and each new suburb has to go to the greatest Australian who does not yet have a suburb. Suppose there are two candidates for the next suburb, and you have to decide between them. [O]n investigating their cases, you conclude it is false that Wye is a greater Australian than Exe, but that it is not false that Exe is a greater Australian than Wye. [...] It would be quite wrong to give the suburb to Wye. *Since the prize was for being the greater Australian, it could not be so obvious who should win unless that person was the greater Australian.*’ [Broome 1997: 74-5, my italics]

As I understand it, Broome’s argument here is an argument to the conclusion that Exe is (determinately) a greater Australian than Wye (and so that it is [determinately] better to award the prize to Exe than to Wye). Its key premise is expressed in the italicised sentence, which I interpret as follows: if you (determinately) ought to award the prize to Exe, then Exe is (determinately) a greater Australian than Wye. Given the new vagueness view, this premise is the culprit in Broome’s argument. To see this, note that, on the new vagueness view, Broome is right that in *Canberra Suburbs* you (determinately) ought to award the prize to Exe, since you ought to believe that awarding the prize to Exe is not worse than awarding it to Wye (because you know that Exe is not less great an Australian than Wye), and you ought not to believe that awarding the prize to Wye is not worse than awarding it to Exe (because you know that Wye is not a greater Australian than Exe). However, according to the new vagueness view, he is wrong that it follows from this that Exe is (determinately) a greater Australian than Wye, since the existence of a requirement to award the prize to Exe is compatible, on the new vagueness view, with another possibility: that it is vague whether Exe is a greater Australian than Wye, and that Wye is (determinately) not a greater Australian than Exe; and so compatible with the possibility that awarding the prize to Exe is (determinately) not worse than awarding it to Wye, and that it is vague whether awarding the prize to Wye is not worse than awarding it to Exe—a case of asymmetric vague betterness.

## 5. THE MAKE NO MISTAKE ARGUMENT, GENERALISED

In Section 3, I developed the Make No Mistake argument against Trichotomy-denying views. As it stands, however, the scope of my argument is restricted: it is conditional on supervaluationism (though this is a standard assumption in the superhard comparisons literature). Is the argument generalisable to other theories of vagueness? In Section 5.1, I provide a simple recipe to generalise it. In Section 5.2, I consider a theory of vagueness to which it cannot be generalised.

### 5.1. *A recipe for generalisation*

There is a simple recipe to generalise the Make No Mistake argument beyond supervaluationism: show that the relevant theory of vagueness entails *Must Not Believe It*, which, recall, states that

*Must Not Believe It*: If one knows that it is vague whether  $p$ , then one ought not to believe  $p$  and one ought not to believe  $\text{not-}p$ .

To see why showing this is enough to generalise the argument, recall the situation in the upper borderline area of B1: agents know that *Operation-UBORDER* is *not worse than* *Not Operating* and that it is vague whether *Not Operating* is *not worse than* *Operation-UBORDER*. Thus, they ought to believe the former proposition and, *by* *Must Not Believe It*, ought not to believe the latter. By Make No Mistake, it follows that they rationally ought to choose *Operation-UBORDER* over *Not Operating*. Therefore, if a theory of vagueness entails *Must Not Believe It*, one can run the Make No Mistake argument, conditional on that theory, and conclude that Trichotomy-denying views are fatally committed to a normative s-configuration without vagueness.

Whether a theory of vagueness entails *Must Not Believe It* is an open question for some theories of vagueness. However, there is a straightforward argument that standard,

Williamsonian epistemicism entails Must Not Believe It. According to epistemicism, *Brad Pitt is tall* is either true or false but we are ignorant of its truth value, since the boundaries of vague predicates like ‘is tall’ are sharp but unknowable [Williamson 1994]. Pairing epistemicism with Williamsonian, knowledge-first epistemology, Must Not Believe It follows. To see this, consider the following thesis of knowledge-first epistemology:

*Knowledge Norm of Belief (KNB)*: One may: believe  $p$  iff one knows that  $p$ .<sup>23</sup>

Given epistemicism, if one knows that it is vague whether *Brad Pitt is tall*, then one knows that one does not know that *Brad Pitt is tall*. If one knows that one does not know that *Brad Pitt is tall*, it follows, by KNB, that one ought not to believe *Brad Pitt is tall* and one ought not to believe its negation. Therefore, together with KNB, epistemicism entails Must Not Believe It.

Can it be shown without appealing to KNB that epistemicism entails Must Not Believe It? I think that this is an open question, but some tentative considerations in favour are given in a footnote.<sup>24</sup>

## 5.2. *A theory of vagueness to which the Make No Mistake argument cannot be generalised*

I have just argued that *any* Must Not Believe It–entailing theory of vagueness is a theory that, if coupled with the Make No Mistake argument, entails that Trichotomy-denying views are false. Supervaluationism and Williamson’s epistemicism are Must Not Believe It–entailing

---

<sup>23</sup> See Williamson [2000] as well as Sutton [2005] and Littlejohn [2013].

<sup>24</sup> If one knows that it is vague whether *Brad Pitt is tall*, then, given epistemicism, how confident ought one to be in *Brad Pitt is tall*? This much seems clear: one ought to have a non-zero degree of confidence (e.g., 0.5), but not full confidence (i.e., 1), in it (and similarly for *Brad Pitt is not tall*, so that the sum of one’s degree of confidence in *Brad Pitt is tall* and *Brad Pitt is not tall* equals 1). Does one’s degree of confidence give one sufficient reason to believe that Brad Pitt is tall? It seems reasonable to think not. After all, one knows that it is *unknowable* whether he is tall. Thus, it seems reasonable to think that this can be shown: given epistemicism, if one knows that it is vague whether  $p$ , then one’s degree of confidence in  $p$  does not give one sufficient reason to believe  $p$ , and one’s degree of confidence in not- $p$  does not give one sufficient reason to believe not- $p$ —and so one ought not to believe  $p$  and ought not to believe not- $p$ . If this can indeed be shown, then it can be shown without appealing to KNB that epistemicism entails Must Not Believe It. For an argument along these lines, see Barnett [2009: 21; MS].

theories. Are there theories that do not entail Must Not Believe It and leave Trichotomy-denying views immune to the argument? Since theories of vagueness are many and varied, it should be unsurprising that there are.

Take, for instance, a theory known as non-standard supervaluationism.<sup>25</sup> On this view, truth is truth on *the correct* sharpening, but it is vague which sharpening is the correct one. If  $p$  is true on all sharpenings of (say) ‘is tall’, then, no matter which is the correct sharpening,  $p$  is *determinately true*. If  $p$  is false on all sharpenings, then, no matter which is the correct sharpening,  $p$  is *determinately false*. If, however,  $p$  is true on some sharpenings but false on others, then it is *vague* (indeterminate) whether  $p$  is true or false, since it is vague whether any of the sharpenings on which it is true is the correct sharpening. *Brad Pitt is tall* is either true or false, but it is vague which.

Consider the implications of this theory for rational belief, on the assumption that TNB holds. If one knows that *Brad Pitt is tall* is true only on some sharpenings and that it is vague whether any of these is the correct sharpening, then, if one believes *Brad Pitt is tall*, it is vague whether one’s belief conforms to TNB. Thus, it is *vague* whether one ought not to believe *Brad Pitt is tall*. To put it another way, non-standard supervaluationism entails not Must Not Believe It but rather:

*Vague Whether One Must Not Believe It*: If one knows that it is vague whether  $p$ , then it is vague whether one ought not to believe  $p$  and vague whether one ought not to believe not- $p$ .

Conditional on non-standard supervaluationism, the situation in the upper borderline area of B1 is one in which agents ought to believe *Operation-UBORDER is not worse than Not Operating*

---

<sup>25</sup> See McGee and McLaughlin [1995] for a defence of it. In more recent years, the basic framework of the theory has been used to motivate the intelligibility of ontic vagueness (and, more generally, to provide a theory of metaphysical indeterminacy). See Barnes [2010] and Barnes and Williams [2011]. For a critique of Barnes and Williams’ approach, see Eklund [2011]. For an alternative way of motivating the intelligibility of ontic vagueness, built around the degree theory of vagueness, see Rosen and Smith [2004].

and, *by* Vague Whether One Must Not Believe It, it is vague whether they ought not to believe *Not Operating is not worse than Operation-UBORDER*. Thus, it is *vague* whether Operation-UBORDER rationally ought to be chosen over Not Operating, which entails that Trichotomy-denying views are committed to a normative s-configuration *with* vagueness.

Non-standard supervaluationism is not without its problems, however.<sup>26</sup> For our purposes, one problem is particularly pressing for any Trichotomy denier who wishes to endorse such a theory to avoid the Make No Mistake argument. Given what it implies for rational choice in cases of asymmetric vague betterness, the theory yields counter-intuitive results in the relevant test cases (Toys, Toys\* and Schock Prize).

At this point, faced with that problem, Trichotomy deniers might want to give up the strategy of *avoiding* the Make No Mistake argument (which consists in endorsing a theory of vagueness to which the argument cannot be generalised, such as non-standard supervaluationism when paired with TNB). In its place, they might feel tempted to adopt a strategy of *rejecting* the Make No Mistake argument such as the following: to endorse (standard) supervaluationism<sup>27</sup> but argue that TNB is false, and thereby reject (rather than merely avoid) the Make No Mistake argument under supervaluationism. The trouble with this strategy is that the argument can go through under supervaluationism even if TNB is false. Why? One obvious reason is that pairing supervaluationism with KNB also entails Must Not Believe It.<sup>28</sup> Moreover, it might well be that supervaluationism entails Must Not Believe It independently of either of those two norms of belief, and so this possibility cannot be ruled out without argument either. (Some reasons for thinking that this is the case are given in a footnote.<sup>29</sup>)

---

<sup>26</sup> See Williamson [1994: 164; 2004: 116-8].

<sup>27</sup> To be clear, whenever I write ‘supervaluationism’ without ‘non-standard’ before it, I always mean to refer to *standard* supervaluationism.

<sup>28</sup> *Proof.* Given supervaluationism, if one knows that it is vague whether  $p$ , then one knows that one cannot know that  $p$  and that one cannot know that not- $p$ , since one knows that knowledge implies truth and that neither  $p$  nor not- $p$  is true. If one knows that one cannot know that  $p$  and that one cannot know that not- $p$ , then one knows that one does not know that  $p$ , and one knows that one does not know that not- $p$ . By KNB, it follows that one ought not to believe  $p$  and ought not to believe not- $p$ .

<sup>29</sup> Given supervaluationism, if one knows that it is vague whether  $p$ , how confident ought one to be in  $p$ ? It seems reasonable to think that one ought not to have any confidence in it and ought not to have any confidence in its

Thus, if Trichotomy deniers felt tempted to adopt a strategy along the lines of the one which we are considering, what they would really need to show is that, despite appearances to the contrary, supervaluationism is not a Must Not Believe It—entailing theory of vagueness.

In any case, the central thesis of this chapter stands: conditional on any Must Not Believe It—entailing theory of vagueness, Trichotomy-denying views are false.

## 6. DECISION-MAKING UNDER VAGUENESS

In previous sections of this chapter, I argued for a new interpretation of the vagueness view of superhard comparisons—which I called the *new* vagueness view to distinguish it from Broome’s (interpretation of the) vagueness view—based on an account of decision-making under vagueness. In this section, I distinguish my account of decision-making under vagueness from other accounts in the literature, namely, those of Robert Williams [2014a, 2016, 2017], Adam Bales [2018] and Susanna Rinard [2015]. The plan is as follows. In Sections 6.1–2, I show that Williams’ and Bales’ accounts are accounts of decision-making under *normative* vagueness (viz., under vagueness in permissibility) while Rinard’s is an account of decision-making under *doxastic* vagueness (viz., under vagueness in belief). In Section 6.3, I argue that these two types of account should be distinguished from a third, namely, from an account of decision-making under *evaluative* vagueness (viz., under vagueness in betterness). Throughout, and in line with the accounts of Williams, Bales and Rinard, vagueness is understood in supervaluationist terms.<sup>30</sup>

---

negation (i.e., one ought to have a zero degree of confidence in  $p$  and a zero degree of confidence in not- $p$ ) [Field 2000; Williams 2017]. After all, one knows that it is not true that  $p$  and that it is not true that not- $p$ . Thus, it seems reasonable to think that supervaluationism entails Must Not Believe It independently of whether TNB holds as well as independently of whether KNB holds, since if one ought not to have any confidence in  $p$  and ought not to have any confidence in not- $p$  when one knows that it is vague whether  $p$ , then one ought not to believe  $p$  and ought not to believe not- $p$ .

<sup>30</sup> A note on terminology: Williams, Bales and Rinard all speak of ‘indeterminacy’ rather than ‘vagueness’, but—to be consistent with the rest of the chapter and because the difference between vagueness and indeterminacy does not matter here (after all, I am understanding vagueness in supervaluationist terms throughout this section)—I prefer to speak of ‘vagueness’. For more on the difference between vagueness and indeterminacy, see Chapter 2.



## 6.1. *Decision-making under normative vagueness*

### 6.1.1. *Williams' account*

Williams' account of decision-making under vagueness is built upon a theory of vagueness that we have just encountered: non-standard supervaluationism [Williams 2017: 666]. Together with TNB, as we saw, non-standard supervaluationism entails:

*Vague Whether One Must Not Believe It:* If one knows that it is vague whether  $p$ , then it is vague whether one ought not to believe  $p$  and vague whether one ought not to believe not- $p$ .

With this in mind, consider the following case (due to Williams [2014a: 2-3] himself):

*Investment:* In a few weeks, Alpha will step into Peter van Inwagen's [1990] Cabinet and, as a result, someone (call him 'Omega') will step out of it. The Cabinet is a machine that changes one's psychology and physiology such that it is vague whether the person who steps into it is the same person as the one who steps out. So, it is vague whether Alpha is Omega. Alpha is aware of all this. Meanwhile, a broker approaches him with an investment opportunity: if Alpha invests a small amount of money now (say, £100), Omega will receive a large amount of money in return. Alpha's preferences are such that if he was Omega, he would invest; and if he was not Omega, he would instead spend the £100 partying before entering the Cabinet. Ought Alpha to invest?

Since Alpha knows that it is vague whether he is Omega, it follows—by Vague Whether One Must Not Believe It—that it is vague whether he ought not to believe the proposition *Alpha is Omega*, and vague whether he ought not to believe the proposition *Alpha is not Omega*. Thus, it is vague whether rationality permits Alpha to invest and vague whether it permits him to not invest, and correspondingly it is vague what he rationally ought to do [Williams 2014a: 25-6]. Given this, Alpha is in a choice situation in which there is normative vagueness.

Williams' question is this: given that there is normative vagueness in his choice situation (i.e., vagueness about what he may and ought to do), what is it that Alpha may and ought to do? Clearly, on pain of incoherence, there must be two senses of (rational) permissibility and requirement at play here: on the first, *strong* sense of permissibility and requirement, there is vagueness about what Alpha may and ought to do; on the second, *weak* sense of permissibility and requirement, there is no vagueness about what Alpha may and ought to do [Williams 2014a: 26; 2016: 423-4]. Accordingly, let us say that there is vagueness about what Alpha is *strongly* permitted and required to do, and no vagueness about what he is *weakly* permitted and required to do. How to make sense of this?

According to Williams' account, while it is vague what Alpha is strongly required to do (i.e., vague whether he is strongly required to invest, and vague whether he is strongly required not to invest), he is weakly required to *make up his mind* how to act—in this case, to invest or to not invest—and it is weakly permissible for him to make up his mind either way. To see this, suppose first that Alpha makes up his mind to, say, invest. Then he does something that is permissible under the sharpening on which Alpha is Omega, and so he does something that is *not* determinately (strongly) impermissible (because it is permissible under at least one sharpening). Now suppose that someone, who knows the choice situation Alpha is in, criticises him by claiming that he acted impermissibly. Then this person is doing something that is only permissible under the sharpening on which Alpha is not Omega. In other words, this person is criticising Alpha while taking a stance (i.e., while not being neutral) on the question of whether Alpha is Omega. So, by investing, Alpha is immune to criticism from a *neutral* standpoint—that is, from the point of view of someone who does not take a stance on the question of whether Alpha is Omega. (*Mutatis mutandis* if we suppose that Alpha makes up his mind to not invest: he would also be immune to criticism from a neutral standpoint.) On Williams' account, it is this immunity to neutral criticism that makes it *weakly* permissible for Alpha to invest (and, *mutatis mutandis*, weakly permissible for him to not invest). More generally,

this is a feature of Williams' account: in any choice situation, choosing an option is weakly permissible iff choosing it is not determinately (strongly) impermissible [Williams 2017: 670].

What is it, however, for one to make up one's mind how to act in the face of normative vagueness? According to Williams [2017: 669-70], it is for one to make an arbitrary, yet binding, judgment call. To unpack this, take a choice situation in which, like Investment,  $\varphi$ ing is (strongly) permissible iff  $p$ , where it is vague whether  $p$ , and assume that one knows that one is in this choice situation. On Williams' account, to make up one's mind how to act in this choice situation is to randomly choose an option— $\varphi$ ing or not  $\varphi$ ing—in a way that constrains future choices. More specifically, to make up one's mind to  $\varphi$  is to adopt a commitment to resolve that normative vagueness, now and on any occasion in the future, in accordance with the sharpening on which  $p$ ; similarly, to make up one's mind to not  $\varphi$  is to adopt a commitment to resolve that normative vagueness, now and on any occasion in the future, in accordance with the sharpening on which not- $p$ .<sup>31</sup> Moreover, to adopt an open-ended commitment to  $\varphi$ —whenever  $\varphi$ ing is (strongly) permissible iff  $p$ , it is vague whether  $p$ , and one knows that one is in this choice situation—is to commit oneself to the judgment that  $p$ ; similarly, to adopt an open-ended commitment to not  $\varphi$  is to commit oneself to the judgment that not- $p$ .

To illustrate this with Investment, suppose that Alpha makes up his mind to (say) not invest. Then he is adopting an open-ended commitment to act under the sharpening on which Alpha is not Omega, which amounts to committing himself to the judgment that Alpha is not Omega. Accordingly, if in the future Alpha faces another choice situation in which  $\varphi$ ing is (strongly) permissible iff Alpha is Omega, he is (rationally) bound to not  $\varphi$ , since he now judges that Alpha is not Omega.

---

<sup>31</sup> Williams argues for this based on considerations having to do with sequences of choices. See Williams [2016: 424-5; 2014: esp. 13-18].

### 6.1.2. Bales' account

As we have seen, a choice situation in which there is normative vagueness is a situation in which it is vague whether it is permissible to choose at least one of the options—call this a *vaguely permissible option*. Are vaguely permissible options ever choiceworthy, and if so when? The aim of Bales [2018: 1693-4] is to answer this question, his answer to it being his own account of decision-making under normative vagueness, according to which vaguely permissible options are sometimes choiceworthy and sometimes not choiceworthy.

To understand Bales' account, consider the following case (from Bales [2018: 1696] himself drawing on Investment):

*Alice's Choice:* Alice will soon step into the Cabinet, following which someone—After-Alice—will step out of it. Before Alice steps into it, though, she has to make a choice between three options. If she chooses the first option ('Wealth'), After-Alice will be paid £20,000. If she chooses the second option ('Pittance'), After-Alice will be paid £1. If she chooses the third option ('Generous'), After-Alice will see Alice's brother, Bob, be paid £100,000. Now, if Alice *is* After-Alice, Alice would prefer to take Generous over Wealth (because she feels that, if she survives and so After-Alice is her, she has a familial duty towards Bob), and she would prefer to take Wealth over Pittance (because she would rather have more money than less). On the other hand, if Alice is *not* After-Alice, Alice would be indifferent between Wealth and Pittance (because she only cares about the money that she and Bob receive), and she would prefer either of those options over Generous (because she actually hates Bob and feels that, if she does not survive and so After-Alice is not her, she has no familial duty towards him). What is it that Alice ought to do?

Bales holds that this is a choice situation, like Investment, in which there is normative vagueness, each of the options—Wealth, Pittance and Generous—being vaguely permissible. However, he also holds that, unlike Investment, there is one option in Alice's Choice that is (determinately) *not* choiceworthy, namely, Pittance. Why? Because Pittance is *vagueness dominated* by Wealth, and no vagueness-dominated option is choiceworthy. According to Bales [2018:

1695], an option,  $y$ , is vagueness dominated by another,  $x$ , if  $x$  is at least as good as  $y$  on all sharpenings and  $x$  is better than  $y$  on at least one sharpening. Vagueness-dominated options (e.g., Pittance) are never choiceworthy because there is no reason to choose them when one could choose the dominating option instead (viz., Wealth).

On Bales' view, then, Alice's Choice is an example of a choice situation in which one of the vaguely permissible options is not choiceworthy. Bales takes this to show at least two things. First, that Williams' account of decision-making under normative vagueness cannot be right, since the account implies that, because it is vague whether Pittance is (strongly) permissible, Pittance is weakly permissible and so choiceworthy. Second, that we need an account of decision-making under normative vagueness on which Pittance, and indeed any vaguely permissible option that is vagueness dominated, is not choiceworthy.

There is more than one account that yields the required result, however. According to what I will call (following Bales) the *modest* account, in any choice situation, if an option is determinately permissible, then it is choiceworthy; if an option is determinately impermissible, then it is not choiceworthy; and a vaguely permissible option is choiceworthy iff it is not vagueness dominated [Bales 2018: 1698]. This account yields the required result because, according to it, *no* vaguely permissible option that is vagueness dominated is choiceworthy. According to what I will call (again, following Bales) the *hedging* account, in any choice situation, an option is choiceworthy iff it is not vagueness dominated [Bales 2018: 1699]. This account yields the required result because, according to it, no option that is vagueness dominated—and so no *vaguely permissible* option that is vagueness dominated—is choiceworthy. Which of these is Bales' account of decision-making under normative vagueness? Ultimately, he remains neutral about this, since he thinks that good cases can be made for both the hedging account and the modest account.<sup>32</sup>

---

<sup>32</sup> Strictly speaking, Bales thinks that, as formulated above, the hedging account faces a problem but that a good case can be made for a reformulated (and more complicated) version of it, one according to which: in any choice situation, an option is choiceworthy iff it maximises *option value* (OV), where an option's OV is calculated by

To see why he thinks this, consider a three-option choice situation in which there is normative vagueness (i.e., in which there is at least one vaguely permissible option), and suppose that there are only two sharpenings. On the first sharpening:  $x > y > z$ . On the second:  $z > y > x$ . Is  $y$  choiceworthy? According to the hedging account,  $y$  is choiceworthy, since it is not vagueness dominated by either  $x$  or  $z$  (on the first sharpening,  $y$  is better than  $z$ ; on the second,  $y$  is better than  $x$ ). On this account, choosing  $y$  is a form of hedging, in this case a way of choosing an option that is *determinately* neither the worst nor the best of the three options. By contrast, the modest account is compatible with  $y$  not being choiceworthy, and in fact there is a *prima facie* compelling argument—available to the proponents of the modest account—to the conclusion that  $y$  is *not* choiceworthy. After all, proponents of the modest account can argue, plausibly, that choosing  $y$  is *determinately* impermissible, since it is impermissible to choose it on the first sharpening (because, on this sharpening,  $x$  is better than  $y$ ) and impermissible to choose it on the second sharpening (because, on this sharpening,  $z$  is better than  $y$ ). If choosing  $y$  is indeed determinately impermissible, then, given the modest account,  $y$  is not choiceworthy.

Regardless of which of these accounts Bales endorses, this much is clear: the key feature of his account of decision-making under normative vagueness is that, in any choice situation, if a vaguely permissible option is vagueness dominated, then it is not choiceworthy.

## 6.2. Decision-making under doxastic vagueness

### 6.2.1. Rinard's account

Rinard's account of decision-making under vagueness is developed as part of a larger project: that of addressing a challenge to the standard model of imprecise credences, which Rinard

---

summing the value of the relevant option on each sharpening [Bales 2018: 1670]. (He calls the reformulated version the *cardinal* hedging account, and the original version the *ordinal* hedging account.) However, for my purposes in what follows, the difference between these two versions of the account does not matter, and so I keep using the original one for simplicity's sake.

calls the *set of functions* model. According to this model, our credence in certain propositions is best represented by a *set* of credence functions, as opposed to *single* credence function (as in standard Bayesianism). To represent my credence in the proposition *I will find a four-leaf clover next Saturday* by a single real number between 0 and 1 would be to represent it with artificial precision—hence the basic motivation for the set of functions model. One challenge that the set of functions model faces, however, is that it is unclear whether there is an adequate account of decision-making to accompany it. In fact, some have argued that there is not.<sup>33</sup> This is the challenge to the set of functions model that Rinard addresses.

What is her strategy to address the challenge? To provide an interpretation of the set of functions model and argue that this interpretation yields an adequate account of decision-making. As becomes clear in what follows, the account that her interpretation of the model yields is (what I call) an account of decision-making under *doxastic* vagueness.

Rinard’s interpretation of the model is built upon non-standard supervaluationism.<sup>34</sup> On this interpretation, every function in an agent’s set of functions is an admissible sharpening of that agent’s epistemic state. Moreover, there is (exactly) one function in the agent’s set such that that is the agent’s credence function, but it is vague which it is. Accordingly, if a proposition  $p$  is true on all functions in the agent’s set, then  $p$  is determinately true of the agent; and if  $p$  is true on some, but not all, functions in the set, then it is vague whether  $p$  is true of the agent [Rinard 2015: 2]. For instance, suppose that  $\text{Pr}(A) > \text{Pr}(B)$  is true on all functions in the agent’s set. Then it is determinately true that the agent has a higher credence in A than in B. If we suppose, however, that  $\text{Pr}(A) > \text{Pr}(B)$  is true on some, but not all, functions in the agent’s set, then it is vague whether it is true that the agent has a higher credence in A than in B.

---

<sup>33</sup> See most notably Elga [2010].

<sup>34</sup> She does not acknowledge this, but it is clear from how she spells out supervaluationism: ‘If [a proposition] is true according to some, but not all, admissible precisifications, then it’s [vague] whether it’s true’ [Rinard 2015: 2]. (Recall that standard supervaluationism entails truth-value gaps, not vague truth.) For discussion of Rinard’s view in connection to non-standard supervaluationism, see Williams [2017: 668-9].

To spell out the account of decision-making that this interpretation yields, notice first that decision-making on the set of functions model comes with a complication that is absent from the single function model. On the latter model, for any choice situation between  $x$  and  $y$ , choosing  $x$  is rationally permissible if it has at least as high expected value as choosing  $y$ , and choosing  $x$  is rationally impermissible if choosing  $y$  has a higher expected value than choosing  $x$ . On the set of functions model, however, there might be choice situations in which choosing  $x$  has at least as high expected value as choosing  $y$  on some functions in the agent's set, and choosing  $y$  has a higher expected value than choosing  $x$  on the other functions in the set. This is a complication because these situations raise the question of which option it is permissible to choose. Rinard's interpretation of the set of functions model offers a novel answer to this question: it is *vague* which option it is permissible to choose.

More specifically, here is the account of decision-making that her interpretation yields. If choosing  $x$  has at least as high expected value as choosing  $y$  on all functions in the agent's set, then choosing  $x$  is permissible on all functions in the set, and so the proposition *Choosing  $x$  is permissible* is determinately true. If choosing  $y$  has a higher expected value than choosing  $x$  on all functions in the agent's set, then choosing  $x$  is *impermissible* on all functions in the set, and so the proposition *Choosing  $x$  is impermissible* is determinately true. Now, if the functions in the agent's set disagree—with some assigning at least as high expected value to choosing  $x$  as to choosing  $y$ , and others assigning a higher expected value to choosing  $y$  than to choosing  $x$ —then choosing  $x$  is permissible on some functions and impermissible on others. So, it is vague whether *Choosing  $x$  is permissible* is true (and thus vague whether *Choosing  $x$  is impermissible* is true).

This is an account of decision-making in which vagueness in what the agent believes (doxastic vagueness) gives rise to vagueness in permissibility (normative vagueness) [Rinard 2015: 3]. To make this clearer, consider, with Rinard [2015: 4], any choice situation in which: an agent is either in belief state BS1 or in BS2, but it is vague which; if the agent is in BS1, then choosing  $x$  is permissible; and if the agent is in BS2, then choosing  $x$  is impermissible. This is



a choice situation in which there is doxastic vagueness. Given Rinard's account, it is vague whether choosing  $x$  is permissible, since the permissibility of choosing  $x$  depends on which belief state, BS1 or BS2, the agent is in, and it is vague in which of these the agent is. Accordingly, if there is doxastic vagueness in a choice situation (as described), there is normative vagueness—this is Rinard's account of decision-making under doxastic vagueness.

### 6.3. *Decision-making under evaluative vagueness*

As should by now be clear, the account of decision-making under vagueness that I developed in Section 3 is neither an account of decision-making under normative vagueness nor an account of decision-making under doxastic vagueness. It is an account of decision-making under *evaluative* vagueness (more specifically, under asymmetric vague betterness). Why does this difference matter? Because failing to distinguish the first two types of account from the third one obscures a significant possibility in logical space: that, conditional on some theory of vagueness, there might be cases of evaluative vagueness—for example, cases of asymmetric vague betterness—that do *not* give rise to normative or doxastic vagueness. In fact, what I did in Section 3 was to develop an account of decision-making under evaluative vagueness that exploits this possibility. After all, there I argued that, conditional on (standard) supervaluationism, the right account of decision-making under asymmetric vague betterness is one according to which there is no normative vagueness (i.e., no vagueness about what the relevant agents may and ought to do) and no doxastic vagueness (i.e., no vagueness about what the relevant agents believe).

To make this clearer, consider first how Toys—a situation in which Pietro faces a choice between Red Toy and Borderline-Red Toy—would be analysed given Rinard's account. Unsurprisingly, it would be analysed as a choice situation in which there is doxastic vagueness: while Pietro (determinately) believes that Red Toy is red, it is vague which belief state, BS1 or BS2, he is in regarding Borderline-Red Toy. If he is in BS1, he believes that Borderline-Red

Toy is red. If he is in BS2, he believes that Borderline-Red Toy is not red. Accordingly, if Pietro is in BS1, it is permissible for him to choose Borderline-Red Toy, in which case he may choose either option; if he is in BS2, it is impermissible for him to choose Borderline-Red Toy, in which case he ought to choose Red Toy. Since it is vague which of the two belief states Pietro is in, it is vague whether it is permissible for him to choose either option, and vague whether he ought to choose Red Toy.

Notice that this contrasts with how Toys is analysed given my account. In light of Must Not Believe It and the fact that Pietro is a rational agent who knows the choice situation he is in, he does not believe the proposition *Borderline-Red Toy is red* nor does he believe its negation (i.e., *Borderline-Red Toy is not red*). In other words, there is *no* vagueness about what Pietro believes.

Consider now how Toys would be analysed given Williams' and Bales' accounts. As expected, it would be analysed as a choice situation in which there is normative vagueness: while it is determinately permissible for Pietro to choose Red Toy, it is vague whether it is permissible for Pietro to choose Borderline-Red Toy. So, it is vague whether it is permissible for Pietro to choose either option, and vague whether he ought to choose Red Toy. Even though Williams' and Bales' accounts would agree on this, they would disagree on what advice to give Pietro. Williams' account—according to which, in any choice situation, choosing an option is weakly permissible iff choosing it is not determinately (strongly) impermissible—would tell Pietro that it is weakly permissible for him to choose Red Toy and weakly permissible for him to choose Borderline-Red Toy. On the other hand, Bales' account—according to which, in any choice situation, if a vaguely permissible option is vagueness dominated, then it is not choiceworthy—would tell Pietro that Borderline-Red Toy, a vaguely permissible option, is not choiceworthy (since it is vagueness dominated by Red Toy).

Once again, notice that this contrasts with how Toys is analysed given my account. In light of Make No Mistake and the fact that Pietro knows the choice situation he is in, he ought

to choose Red Toy. In other words, there is *no* vagueness about Pietro may and ought to do. Thus, I conclude, accounts of decision-making under normative vagueness, like Williams' and Bales', and accounts of decision-making under doxastic vagueness, like Rinard's, should be distinguished from accounts of decision-making under evaluative vagueness, like mine.

## 7. CONCLUSION

My aim in this chapter was to build the basic framework for the new vagueness view of superhard comparisons and choices. As I initially developed it, one of its building blocks was supervaluationism. I then showed that the Make No Mistake argument, upon which the new vagueness view is built, is generalisable to all Must Not Believe It—entailing theories of vagueness. At the core of the view lies a distinction between symmetric and asymmetric vague betterness, with an account of decision-making under the latter being put forward. If the view is correct, the phenomenon of superhard comparisons and choices is not grounded in incomparability or parity, as Trichotomy-denying views hold. It is grounded in a very familiar phenomenon—vagueness.

## *Chapter 2*

# Superhard choices and moral vagueness

### 1. INTRODUCTION

Consider the following cases:

*Diversions:* You are watching your five-year-old daughter play in the city park. It is morally permissible for you to divert your attention from her for one second. It is not morally permissible for you to divert your attention for six hundred seconds (ten minutes). How many seconds does it take for it to become morally impermissible for you to divert your attention?

*Lives:* You can only save your friend or some strangers. It is morally permissible for you to save your friend rather than a few strangers. It is not morally permissible for you to save your friend rather than very many strangers. How many strangers does it take for it to become morally impermissible for you to save your friend?<sup>35</sup>

---

<sup>35</sup> The first case is adapted from Schoenfield [2016: 262]; the other is adapted from Dougherty [2014: 353].

In each of these cases, we can create a sorites paradox for the predicate ‘is morally permissible’.

Take, for instance, the Diversions case. In Diversions, it seems clear that, if it is morally permissible for you to divert your attention for one second, it is morally permissible for you to divert your attention for two seconds. In fact, the following principle seems to hold:

*Tolerance*<sup>(D)</sup>: If it is morally permissible for you to divert your attention from your daughter for  $n$  number of seconds, then it is morally permissible for you to divert your attention for  $n + 1$  (i.e., for another second).

By *modus ponens* and continuous application of *Tolerance*<sup>(D)</sup>, we can create a *moral* sorites paradox:

*Moral sorites paradox*

- (1) It is morally permissible for you to divert your attention from your daughter for one second.
- (2) If it is morally permissible for you to divert your attention for one second, then it is morally permissible for you to divert your attention for one second + 1.
- (3) If it is morally permissible for you to divert your attention for two seconds, then it is morally permissible for you to divert your attention for two seconds + 1.
- ...
- Therefore,
- (C) It is morally permissible for you to divert your attention from your daughter for six hundred seconds (ten minutes).

That we can create a sorites paradox for ‘is morally permissible’ is strong evidence that there are cases in which it is *vague* whether it is morally permissible for one to  $\varphi$ . In other words, it is evidence that there is *moral vagueness*. In recent years, interest has grown in the implications

of vagueness for morality.<sup>36</sup> How does this relate, however, to the growing interest in superhard comparisons and choices, the topic of Chapter 1?

According to the new vagueness view of superhard comparisons and choices, vagueness in ‘better than’—what I have called *vague betterness*—can arise in two ways: whenever it is vague whether  $x$  is better than  $y$ , and  $y$  is (determinately) not better than  $x$ , what I have called *asymmetric vague betterness*; and whenever it is vague whether  $x$  is better than  $y$  and vague whether  $y$  is better than  $x$ , what I have called *symmetric vague betterness*. In Chapter 1, I developed an account of decision-making under asymmetric vague betterness, according to which in any case of asymmetric vague betterness one ought to choose  $x$  over  $y$ , on pain of being practically irrational. Based on this account, I put forward the Make No Mistake argument against Trichotomy-denying views and built the basic framework for the new vagueness view. As it stands, however, my defence of the new vagueness view is incomplete. After all, in Chapter 1, I did not develop any account of decision-making under symmetric vague betterness. All I did in Chapter 1 was to show which account is needed to fully defend the new vagueness view, namely, an account according to which, in at least some cases of symmetric vague betterness, it is vague which of  $x$  and  $y$  one rationally may choose and which of  $x$  and  $y$  one rationally ought to choose. In a nutshell, I did not show that there is vagueness in rational choice in at least some cases of symmetric vague betterness.

In this chapter, I explore uncharted territory. I consider the problem of decision-making under symmetric vague betterness through the lens of cases of symmetric vague betterness that have their source in moral vagueness. To anticipate, these are case of symmetric vagueness in ‘(all-things-considered) better than’ that have their source in cases of vagueness in ‘is morally permissible’. My aim is twofold. On the one hand, I aim to show that two dominant types of theory of vagueness—which I will call (following David Barnett [2011])

---

<sup>36</sup> See especially Dougherty [2014], Constantinescu [2014], Schoenfield [2016], Dunaway [2017], and Williams [2017].

*vagueness-as-ignorance* and *vagueness-as-absence-of-truth*, the latter of which supervenience is a version—deliver different accounts of decision-making under those cases, and that this has interesting implications for (at least) two debates in practical ethics: the abortion debate and the animal welfare debate. On the other hand, I aim to show that, conditional on *vagueness-as-absence-of-truth*, there is vagueness-induced indeterminacy in rational choice in those cases—just what is needed to complete my defence of the new vagueness view.

The structure of the chapter is as follows. In Section 2, I introduce two cases of moral vagueness—one from the abortion debate and another from the animal welfare debate. In Section 3, I show that there are cases of symmetric vague betterness that have their source in those cases of moral vagueness. In Section 4, I argue that *vagueness-as-ignorance* delivers one account of decision-making under those cases of symmetric vague betterness, and that *vagueness-as-absence-of-truth* delivers another account of decision-making under the same cases. In Section 5, I discuss the implications of these accounts for the abortion and animal welfare debates. In Section 6, I revisit the new vagueness view in light of these results.

## 2. TWO CASES OF MORAL VAGUENESS FROM PRACTICAL ETHICS

### 2.1. *A case of moral vagueness from the abortion debate: Abortions*

Consider the following case:

*Abortions*: Francesca is pregnant but wants to abort because having the child would interfere with her career plans. (Call this *Francesca's reason*.<sup>37</sup>) It is morally permissible for Francesca to abort in the first few days of her pregnancy for her reason. It is not morally permissible for Francesca to abort in the last few days of her pregnancy for her reason. Between a possible abortion carried out in the first few days of the pregnancy and one carried out in the last few days, there is a series, any abortion in it ( $A_i$ ) differing from the one that follows it ( $A_{i+1}$ ) in that  $A_i$  is carried out one millisecond (or nanosecond, or picosecond) earlier than  $A_{i+1}$ —a

---

<sup>37</sup> Francesca's reason is among the most common reasons why women abort. For the relevant empirical data, see Finer et al. [2005].

*sorites* series. At what point in the series does it become morally impermissible for Francesca to abort for her reason?<sup>38</sup>

To generate a moral sorites paradox on the basis of Abortions, we need to assume the *liberal view* of the morality of abortion. To that end, I introduce below (§2.1.1) a simple version of the liberal view, endorsed among others by L. W. Sumner [1983, 1985], Bonnie Steinbock [2011], and the legal scholars Sherry F. Colb and Michael C. Dorf [2016].<sup>39</sup> Then I show (§2.1.2) that this view allows us to generate a moral sorites paradox on the basis of Abortions and thereby conclude that there are cases in the series in which it is vague whether it is morally permissible for Francesca to abort for her reason.

#### 2.1.1. *The simple liberal view*

To introduce the simple version of the liberal view, it is helpful to first contrast the basic tenets of any version of the liberal view with the basic tenets of any version of its two rival views. According to the *conservative view*, it is never—or, on some versions, almost never—morally permissible for a pregnant woman to abort, and this is so regardless of the stage of the pregnancy at which she is. For example, if Francesca were at the early stage of her pregnancy (say, the second month) and wanted to abort for her reason (i.e., that having the child would interfere with her career plans), it would be morally impermissible for her to do so.<sup>40</sup> According to (what might be called) the *very liberal view*, it is always—or, on some versions, almost always—morally permissible for a pregnant woman to abort, and this is so regardless of the stage of the

---

<sup>38</sup> Note that, throughout this chapter (and dissertation), I do not rely on any distinction between subjective and objective senses of ‘morally permissible’—a distinction according to which there is, respectively, a sense of ‘morally permissible’ that is sensitive to the epistemic circumstance in which agents happen to be at the time of action (what their beliefs are and what their evidence is in that circumstance), and a sense of ‘morally permissible’ that is sensitive to the underlying facts, regardless of the epistemic circumstance in which agents happen to be. Those who believe in such a distinction should interpret my use of ‘morally permissible’ as referring to the objective sense of ‘morally permissible’.

<sup>39</sup> Presumably, there are also views in the literature on abortion that are compatible, but do not entail, the simple version of the liberal view. I suspect that Elizabeth Harman’s [1999] is an example of such a view.

<sup>40</sup> For an influential defence of the conservative view, see Marquis [1989]. For more recent defences, see Lee [2010] and especially Kaczor [2015].



pregnancy at which she is. To keep using our example, if Francesca were at the late stage of her pregnancy (say, the eighth month) and wanted to abort for her reason, it would be morally permissible for her to do so.<sup>41</sup>

The liberal view contrasts with both of those views in holding that it is always morally permissible for a pregnant woman to abort at the early stage of her pregnancy and almost never morally permissible for her to do so at the late stage of it. Thus, if Francesca were at the early stage of her pregnancy and wanted to abort for her reason, it would be morally permissible for her to do so; whereas if she were at the late stage of the pregnancy and wanted to abort for the same reason, it would be morally impermissible for her to do so.<sup>42</sup>

A simple (and common) way of defending the liberal view goes as follows. Sentience, the capacity that certain beings have to feel pain or pleasure, is necessary and sufficient for having interests—crucially, for having an interest in not being inflicted pain and an interest in not being deprived of a life worth living. Given this, foetuses have interests iff they are sentient. This implies that, when deliberating whether to abort, a woman who is pregnant with a *non-sentient* foetus should not consider its interests (because it has none).<sup>43</sup> Accordingly, it is morally permissible for her to abort for any reason.<sup>44</sup> As for a woman who is pregnant with a

---

<sup>41</sup> For influential defences of the very liberal view, see Warren [1973], Tooley [1972] and Singer [2011: 123-154]. For a more recent defence, see especially McMahan [2002: 267-422].

<sup>42</sup> For an influential defence of the liberal view, see Thomson [1971]. If Rosalind Hursthouse's [1991] virtue theoretic approach to the morality of abortion is best understood as a version of the liberal view, as I suspect is the case, then hers is another influential defence of the view. For other defences, see especially Catriona Mackenzie's [1992]—which is partially a critique of both the liberal view of Judith Jarvis Thomson [1971] and the very liberal view of Mary Anne Warren [1973]—as well as Margaret Little's [2008] and Christine Overall's [1987: 40-87] (though Overall has recently repudiated her version of the liberal view in Overall [2015], where she endorses a version of it more in line with Mackenzie's).

<sup>43</sup> See, for example, Steinbock [2011: 58]: 'A late-gestation fetus is conscious and sentient. It is very likely that it can experience pain, and presumably, it has pleasurable experiences. If so, it has an interest in continuing to live and in having those pleasurable experiences, an interest that can be protected by a right to life. By contrast, embryos and preconscious fetuses do not have lives that they value, lives that are a good to them, because they are unaware of everything around them. [...] Thus, the importance of sentience is not primarily that abortion causes pain to the sentient fetus. That problem might be taken care of with an anesthetic. The relevance of sentience is that a sentient being can have a life it values, and that we can protect for its own sake.'

<sup>44</sup> See, for example, Sumner [1983: 152]: 'The moral issues raised by early abortion [i.e., abortion of non-sentient foetuses] are precisely those raised by contraception. [...] Since the fetus at this stage has no right to life, early abortion (like contraception) cannot violate its rights. But if it violates no one's rights, early abortion (like contraception) is a private act. There are of course significant differences between contraception and early abortion, since the former is generally less hazardous, less arduous, and less expensive. A woman has, therefore,

*sentient* foetus, it is only morally permissible for her to abort for a very strong reason, such as her life being at risk if the pregnancy continues, a health condition with the foetus so severe that its life is not worth living, and a few other reasons. Otherwise, if the reason is not strong enough, it is morally impermissible—and indeed very wrong—for the woman to abort.<sup>45</sup> Since foetuses that do become sentient only become so at some point in the second trimester of the pregnancy<sup>46</sup>, it is always morally permissible for a pregnant woman to abort *before* that point (i.e., to carry out an early-stage abortion) and almost never morally permissible for her to do so *after* that point (i.e., to carry out a late-stage abortion). This is the *simple (version of the) liberal view*.

### 2.1.2. *A moral sorites paradox on the basis of Abortions*

Given the simple liberal view—which I shall assume from now on—it is easy to see how to generate a moral sorites paradox on the basis of Abortions. On the simple liberal view, it is morally permissible for Francesca to abort for her reason at  $A_1$  (the first abortion case in the series, corresponding to, say, a day in the first week of the pregnancy), since the foetus is not sentient. Also on the simple liberal view, it is not morally permissible for Francesca to abort for her reason at  $A_n$  (the last case in the series, corresponding to, say, the day before the birth of the child), since the foetus is sentient and only a very strong reason would make it permissible for her to abort. Is there a sharp boundary between the cases in which it is morally permissible for Francesca to abort for her reason and the cases in which it is not morally permissible for her to do so? It seems not, for it seems that if Francesca acted permissibly by aborting for her reason at  $A_i$  (an arbitrary case in the series), then she would be acting

---

good prudential reasons for relying on contraception as her primary means of birth control. But if she elects an early abortion, then, whatever the circumstances and whatever her reasons, she does nothing immoral.’

<sup>45</sup> See, for example, Colb and Dorf [2016: 13]: ‘Absent some very strong justification or excuse, we have a duty at least to avoid intentionally inflicting suffering or death on any sentient being, whether human or nonhuman.’

<sup>46</sup> See Sumner [1983: 147-150] and Steinbock [46-50]. For a more detailed discussion of when sentience emerges in foetal development, see Boonin [2003: 98-129].

permissibly by aborting one millisecond later for her reason. To put it another way, it seems that the relevant version of Tolerance holds here:

*Tolerance*<sup>(A)</sup>: If it is morally permissible for Francesca to abort for her reason at  $A_i$ , then it is morally permissible for her to abort for her reason at  $A_i + 1$  (i.e., in the next case/one millisecond later).

By *modus ponens* and continuous application of *Tolerance*<sup>(A)</sup>, we can generate the relevant moral sorites paradox:

*Moral sorites paradox on the basis of Abortions*

- (1) It is morally permissible for Francesca to abort for her reason at  $A_1$ .
- (2) If it is morally permissible for Francesca to abort for her reason at  $A_1$ , then it is morally permissible for her to abort for her reason at  $A_1 + 1$ .
- (3) If it is morally permissible for Francesca to abort for her reason at  $A_2$ , then it is morally permissible for her to abort for her reason at  $A_2 + 1$ .
- ...
- Therefore,
- (C) It is morally permissible for Francesca to abort for her reason at  $A_n$ .

In other words, we can generate an apparently sound argument to the conclusion that it is morally permissible for Francesca to abort for her reason the day before the birth of the child ( $A_n$ )—a false conclusion, given the simple liberal view. This is strong reason to believe that there is (what I call) a *zone of vagueness* in the series of Abortions: a zone in which it is vague whether it is morally permissible for Francesca to abort in it for her reason—a zone between the one in which it is *determinately* morally permissible for her to abort in it for her reason and the one in which it is *determinately not* morally permissible for her to abort in it for her reason.<sup>47</sup>

Here is a schematic representation of the series:

---

<sup>47</sup> Interestingly, Sumner is aware of this and in fact embraces the existence of the zone of vagueness as a feature of the simple liberal view (for cases like Abortions). See Sumner [1983: 150; 1985: 695-6].

$A_1, A_2, \dots$

*Zone of vagueness*

$\dots, A_{n-1}, A_n$

---

#### SERIES OF ABORTIONS

When Francesca is in the zone of vagueness of the series of Abortions, her case is a case of moral vagueness.

#### 2.2. *A case of moral vagueness from the animal welfare debate: Meals*

Consider the following case:

*Meals*: Tomaso wants to eat a particular meal because he likes the taste of it and finds it tastier than any of the other options on this new app that he just downloaded, which allows customers to order an astoundingly wide range of meals. (Call this *Tomaso's reason*.) The number and variety of meals on the app is so high that it can be arranged as a sorites series, any meal in it ( $M_i$ ) differing from the one that follows it ( $M_{i+1}$ ) in that  $M_i$  is made from an edible thing that, when alive, had marginally simpler information-processing capabilities than the one from which  $M_{i+1}$  is made. (Assume that, in the process of being reared or killed, the things in the series that could suffer—because they were sentient—did suffer.) It is morally permissible for Tomaso to eat  $M_1$ —the first meal in the series (say, a plant-based meal)—for his reason. It is not morally permissible for Tomaso to eat  $M_n$ —the last meal in the series (say, a pork-based meal)—for his reason. At what point in the series does it become morally impermissible for him to eat a meal for his reason?

To generate a moral sorites paradox on the basis of Meals, we need to assume *moral vegetarianism*. A simple version of moral vegetarianism—which I call *simple moral vegetarianism*—is Peter Singer's [2002, 2011], by far the most influential version of the view.<sup>48</sup> In what follows,

---

<sup>48</sup> Other versions of moral vegetarianism include Tom Regan's [1983] animal rights version, Hursthouse's [2000] virtue theoretic version, and Christine Korsgaard's [2018] Kantian version.

I introduce simple moral vegetarianism (§2.2.1) and then show that it allows us to generate a moral sorites paradox on the basis of Meals (§2.2.2).

### 2.2.1. *Simple moral vegetarianism*

According to simple moral vegetarianism, the basic principle of equality implies that our standard meat production and consumption practices are *speciesist*—that is, objectionable in a similar way to that in which racist or sexist practices are objectionable. The principle in question is the *principle of equal consideration of interests*, which demands us to give equal weight to the interests of others when deliberating what to do. To see the explanatory work that the principle does, consider how racists violate it: by giving more weight to the interests of members of their own race than to the interests of members of other races whenever those interests conflict—hence giving *unequal* weight to the interests of others. Consider also how sexists violate it: by giving more weight to the interests of members of their own sex than to the interests of members of the opposite sex whenever those interests conflict. On simple view vegetarianism, to give more weight to the interests of members of one’s own species than to the interests of members of other species whenever those interests conflict is similarly to violate the principle of equal consideration of interests. Why?

To see why, note that, on simple moral vegetarianism, if something has interests, then it has the capacity to feel pain or pleasure; and if something has the capacity to feel pain or pleasure, then it has interests. In other words, being *sentient* is necessary and sufficient for having interests.<sup>49</sup> Thus, if something is sentient, it must have a basic interest in not suffering, in not being inflicted pain. Why is sentience the boundary that divides that which has interests from that which does not have interests? Because, according to simple moral vegetarianism,

---

<sup>49</sup> In Singer’s [2002: 8] own words: ‘The capacity for suffering and enjoyment is, however, not only necessary, but also sufficient for us to say that a being has interests—at an absolute minimum, an interest in not suffering.’

any way of drawing the boundary through the class of sentient things, such as skin colour or sex, would be arbitrary and violate the principle of equal consideration of interests.

One upshot of this account of interests is that plants, for instance, do not have interests because they do not meet the necessary condition for having them—being sentient.<sup>50</sup> Another is that members of many species other than *Homo sapiens* have interests because they meet the sufficient condition for having them—being sentient. Thus, to give more weight to the interests of members of *Homo sapiens* than to the interests of members of, say, *Sus scrofa* (the species to which domestic pigs belong) whenever those interests conflict is to violate the principle of equal consideration of interests.

A standard case in which the interests of *Homo sapiens* conflict with the interests of *Sus scrofa* is, of course, when some of us—persons who are members of *Homo sapiens*<sup>51</sup>—have an interest in eating members of *Sus scrofa* but they have an interest in not suffering. On simple moral vegetarianism, if that interest some of us have is a *major* interest, such as when the suffering is brought about for us to not die of starvation, then it outweighs the interest of the relevant members of *Sus scrofa* (for even if they are not only sentient but also self-aware, they are still self-aware to a much lesser extent than human persons are) [Singer 2011: 122]. In that case, it is morally permissible to eat the relevant members of *Sus scrofa*. But in standard cases of animal welfare our interest is a *minor* interest, such as when the suffering is brought about in order for us to enjoy gustatory pleasure. Accordingly, if the relevant members of *Sus scrofa* suffer in the process of being reared or killed—as they do in standard cases—that minor interest of ours is outweighed by their interest in not suffering. Thus, at least in standard cases, it is not morally permissible to eat pork and *mutatis mutandis* other kinds of meat, on pain of being speciesist.

---

<sup>50</sup> For an argument that plants are not sentient, see Tye [2017: 168-72].

<sup>51</sup> Singer distinguishes human beings (i.e., members of *Homo sapiens*) from persons.

### 2.2.2. *A moral sorites paradox on the basis of Meals*

Let us now return to Meals. Given simple moral vegetarianism—which I shall assume from now on—while it is morally permissible for Tomaso to eat meal  $M_1$  (a plant-based meal) for his reason (i.e., that he likes the taste of it and finds it tastier than any of the other options on the app), it is not morally permissible for him to eat  $M_n$  (a pork-based meal) for his reason. After all, on simple moral vegetarianism, for Tomaso to eat  $M_n$  for his reason is for him to be speciesist.<sup>52</sup> Is there a sharp boundary between the cases in which it is morally permissible for Tomaso to eat a meal for his reason and the cases in which it is not morally permissible for him to do so? It seems not, for it seems that

*Tolerance*<sup>(M)</sup>: If it is morally permissible for Tomaso to eat  $M_i$  (an arbitrary meal in the series) for his reason, then it is morally permissible for him to eat  $M_i + 1$  (i.e., the next meal in the series) for his reason.

By *modus ponens* and continuous application of *Tolerance*<sup>(M)</sup>, we can generate the relevant moral sorites paradox:


#### *Moral sorites paradox on the basis of Meals*

- (1) It is morally permissible for Tomaso to eat  $M_1$  for his reason.
  - (2) If it is morally permissible for Tomaso to eat  $M_1$  for his reason, then it is morally permissible for Tomaso to eat  $M_1 + 1$  for his reason.
  - (3) If it is morally permissible for Tomaso to eat to eat  $M_2$  for his reason, then it is morally permissible for Tomaso to eat  $M_2 + 1$  for his reason.
  - ...
- Therefore
- (C) It is morally permissible for Tomaso to eat  $M_n$  for his reason.

---

<sup>52</sup> Accordingly, by assuming simple moral vegetarianism, I restrict my discussion to moral vegetarianism as a view that is only concerned with the welfare of animals (and so set aside discussing moral vegetarianism as a view that is also concerned with the environment).

The conclusion of this argument is obviously unacceptable given simple moral vegetarianism. The upshot is that there must be cases in the series of Meals such that it is *vague* whether it is morally permissible for Tomaso to eat the relevant meal for his reason—in addition to cases in which it is *determinately* morally permissible for him to eat the relevant meal for his reason (viz.,  $M_1, M_2, \dots$ ) and in addition to cases in which it is *determinately not* morally permissible for him to eat the relevant meal for his reason (viz.,  $M_n, M_n - 1, \dots$ ). Let us refer to those cases as cases in the *zone of vagueness* of the series of Meals, which might schematically be represented as follows:



When Tomaso wants to eat a meal that is in the zone of vagueness of the series of Meals, his case is a case of moral vagueness.

### 2.3. *A common source*

One feature of our two cases of moral vagueness is that the source of the moral vagueness is the same. To see this, recall that, according to both the simple liberal view and simple moral vegetarianism, being sentient is necessary and sufficient for having interests. In  $A_1$ , the foetus does not have interests because it is not sentient; in  $A_n$ , the foetus has interests because it is sentient. Similarly, in  $M_1$ , the plant (when alive) did not have interests because it was not sentient; in  $M_n$ , the pig (when alive) had interests because it was sentient. In the zone of vagueness of the series of Abortions/Meals, there is vagueness about whether it is morally permissible to abort/eat the relevant meal because it is vague whether the relevant being has/had interests, and thereby vague whether the relevant being has/had sentience. Thus, both



the cases of moral vagueness in the series of Abortions and the cases of moral vagueness in the series of Meals have their source in *vagueness about whether the relevant being has/had sentience*.<sup>53</sup>

One might object that it is controversial to assume, as I do, that there is vagueness about whether certain beings are sentient, since some philosophers of mind think that sentience—and (phenomenal) consciousness more generally—is an on-or-off phenomenon and so that there is always a fact of the matter about whether a being is sentient.<sup>54</sup>

Notice, however, that this view in the philosophy of mind is compatible with there being vagueness about whether certain beings are sentient, since the view does not rule out the possibility that vagueness in sentience be explained in purely *epistemic* terms. Moreover, it should be noted that other philosophers of mind disagree with the ‘on-or-off’ view of sentience, holding that sometimes there is vagueness about whether a being is sentient and, crucially, that this vagueness is to be explained in purely *non-epistemic* terms.<sup>55</sup> What matters for my purposes in what follows then is that, on the assumption that there is vagueness in sentience, it is very natural to pair the ‘on-or-off’ view of sentience with vagueness-as-ignorance and to pair the rival view in the philosophy of mind with vagueness-as-absence-of-truth.<sup>56</sup>

#### 2.4. *A common structure*

Another feature of our two cases of moral vagueness is that they share the same structure. Let ‘ $\phi$ ing’ stand for acts like ‘aborting’ or ‘eating the relevant meal’. In both cases of moral

---

<sup>53</sup> The existence of vagueness in sentience relates more generally to a hotly debated question in the philosophy and science of consciousness: which animals are sentient? Are fish sentient? What about shellfish? For an argument against and another for the view that fish are sentient, see Rose et al. [2014] and Braithwaite [2010], respectively. For the view that sentience is very widespread across the animal kingdom, see Tye [2017]. For a literature review on animal sentience, see Allen and Trestman [2016].

<sup>54</sup> They include Colin McGinn [1982: 13–4], John Searle [1992: 83] and David Chalmers [1996: 105], but the only sustained defences of the view in the literature so far are in Antony [2006a, 2006b], Simon [2017] and O’Rourke [MS].

<sup>55</sup> They include Ned Block [1992: 205; 2002: 418–9], David Papineau [1993: 128–9], Daniel Dennett [1995: 706] and Michael Tye [2000: 180–1; 1996: 682–3].

<sup>56</sup> On vagueness as an epistemic phenomenon, see Section 4.2; on vagueness as a non-epistemic phenomenon, see Section 4.4.

vagueness, as we saw, it is vague whether  $\varphi$ ing is morally permissible. However, notice that, also in both cases, *not*  $\varphi$ ing—not aborting, not eating the relevant meal—is (determinately) morally permissible. In other words, there is no vagueness about whether carrying the pregnancy to term is morally permissible for Francesca and no vagueness about whether refraining from eating the relevant meal is morally permissible for Tomaso; in both cases, it is morally permissible to not  $\varphi$ . Why? On the one hand, because, on any version of the liberal view, Francesca would be doing nothing wrong, morally, by carrying the pregnancy to term, regardless of the stage of the pregnancy at which she is. On the other hand, because, on any version of moral vegetarianism, Tomaso would be doing nothing wrong, morally, by refraining from eating a meal in the zone of vagueness and instead ordering a meal that is (say) plant-based.

### 3. MORAL VAGUENESS AND SYMMETRIC VAGUE BETTERTNESS

In Section 2, I introduced two cases of moral vagueness with the same structure. In this section, I show that there are cases of symmetric vague betterness that have their source in those cases of moral vagueness. To that end, let me sharpen the details of the two cases. Francesca wants to abort but she is in the zone of vagueness of the series of Abortions. More specifically, Francesca’s choice situation is such that it is *prudentially better* for her—that is, better from the point of view of her own interests—to abort than to not abort, the prudential reason which settles that it is prudentially better to abort being that *having the child would interfere with her career plans*.<sup>57</sup> At the same time, it is *vague* whether it is *morally permissible* for her to abort for her reason (where this reason is the prudential reason just mentioned). Call this *Francesca’s moral vagueness case*. Similarly, Tomaso wants to eat a particular meal but that meal is in the zone of vagueness of the series of Meals. More specifically, Tomaso’s choice situation is such that it is

---

<sup>57</sup> I take a prudential reason to be any consideration that counts in favour of an action’s being prudentially good or prudentially bad to do (where ‘action’ throughout this chapter should always be understood broadly, to also include omissions).

*prudentially better* for him to eat that meal than to not eat it, the prudential reason which settles that it is prudentially better to eat the meal being that *he likes the taste of it and finds it tastier than any of the other options on the app*. At the same time, it is *vague* whether it is *morally permissible* for him to eat that meal for his reason (where this reason is the prudential reason just mentioned). Call this *Tomaso's moral vagueness case*.

In both Francesca's and Tomaso's moral vagueness cases, the situation is such that it is prudentially better for a subject S to  $\varphi$  (abort, eat a meal) but it is vague whether it is morally permissible for S to  $\varphi$ . Given this, we can ask the following question:

*Choice Under Moral Vagueness Question:*

What is it that one rationally ought to do if one knows that it is (determinately) prudentially better for one to  $\varphi$ , but that it is vague whether it is morally permissible for one to  $\varphi$ ?

This question asks for something unexplored in the literature: an account of decision-making under such cases. As I argue in what follows, at least some of these cases—namely, Francesca's and Tomaso's moral vagueness cases—are a (proper) subset of cases of symmetric vague betterness. Thus, answering the Choice Under Moral Vagueness Question for Francesca's and Tomaso's cases requires developing a partial account of decision-making under symmetric vague betterness. I take up the task of answering (a more precise version of) the Choice Under Moral Vagueness Question for Francesca's and Tomaso's cases in Section 4. For now, why is it that Francesca's and Tomaso's moral vagueness cases are cases of symmetric vague betterness?

To show that Francesca's and Tomaso's moral vagueness cases are cases of symmetric vague betterness, I show that morality overrides prudence whenever there is a conflict between them in cases in the series of Abortions and in cases in the series of Meals, and that it is better to be prudent whenever there is no conflict between morality and prudence in cases in the series of Abortions and in cases in the series of Meals. In Section 3.1, I discuss the cases in the

series of Abortions and show that Francesca's case is a case of symmetric vague betterness. In Section 3.2, I discuss the cases in series of Meals and show that Tomaso's case is a case of symmetric vague betterness.

### 3.1. *Francesca's moral vagueness case and symmetric vague betterness*

#### 3.1.1. *Late Stage*

Suppose that Francesca is, and knows that she is, at a point in the series of Abortions such that it is *not* morally permissible for her to abort for her reason (i.e., that having a child would interfere with her career plans), but that it is prudentially better for her to abort. (Suppose also, here and throughout, that she knows that the simple liberal view is true.) Call any such case in the series *Late Stage*.

In Late Stage, morality and prudence conflict. There are several views about how morality and prudence interact in cases of conflict.<sup>58</sup> However, I take it that, given the simple liberal view (which, recall, we are assuming), the following claim about the morality/prudence interaction holds for cases in the series of Abortions:

*Morality Always Overrides Prudence In Cases Of Conflict In Abortions:* For any case in the series of Abortions, if it is prudentially better for Francesca to abort but not morally permissible to do so, then not aborting is (all-things-considered) better than aborting.

To see why, suppose that this claim is false. In other words, suppose that

*Morality Sometimes Does Not Override Prudence In Cases of Conflict In Abortions:* For some case in the series of Abortions, if it is prudentially better for Francesca to abort but not morally permissible to do so, then aborting is *either* (all-things-considered) better than *or* (all-things-considered) equally as good as not aborting.

---

<sup>58</sup> One being, of course, Sidgwick's [1962] classic dualism of practical reason. For good recent discussions of this, see Dorsey [2013] and Parfit [2011: 130-49].

Suppose further that Francesca is, and knows that she is, at the point in the series for which this case holds, which entails that she knows that aborting is either (all-things-considered) better than or (all-things-considered) equally as good as not aborting. Then practical rationality *permits* her to abort, even though she knows that at that point in the series the foetus is sentient. I take this to be a *reductio* of Morality Sometimes Does Not Override Prudence In Cases of Conflict In Abortions. As a result, I take it that, given the simple liberal view, Morality Always Overrides Prudence In Cases Of Conflict In Abortions holds. The upshot is that, in Late Stage, *not* aborting is (all-things-considered) *better than* aborting.

### 3.1.2. *Early Stage*

Suppose that Francesca is, and knows that she is, at a point in the series such that it *is* morally permissible for her to abort for her reason, and that it is prudentially better for her to do so. Call such any such case in the series *Early Stage*.

In Early Stage, there is no conflict between morality and prudence. Given the simple liberal view, I take it that the following claim holds for cases in the series of Abortions:

*Always Better Be Prudent In Cases Of No Conflict In Abortions:* For any case in the series of Abortions, if it is prudentially better for Francesca to abort and morally permissible to do so, then aborting is (all-things-considered) better than not aborting.

The upshot is that, in Early Stage, aborting is (all-things-considered) *better than* not aborting.

### 3.1.3. *Francesca's is a case of symmetric vague betterness*

The upshot of these two cases outside the zone of vagueness of the series of Abortions is then this. If it is prudentially better for Francesca to abort but morality does not permit doing so (as in Late Stage), then *not* aborting is (all-things-considered) *better than* aborting. If it is

prudentially better for Francesca to abort and morality permits doing so (as in Early Stage), then *aborting* is (all-things-considered) *better than* not aborting. Accordingly, if it is prudentially better for Francesca to abort but *vague* whether morality permits doing so—as in Francesca’s moral vagueness case—then it is vague whether aborting is (all-things-considered) better than not aborting, *and* vague whether not aborting is (all-things-considered) better than aborting. Francesca’s case is a case of symmetric vague betterness.

### 3.2. *Tomaso’s moral vagueness case and symmetric vague betterness*

#### 3.2.1. *Chain’s End*

Suppose that Tomaso wants to eat a meal in the series of Meals which he knows to *not* be morally permissible to eat for his reason (i.e., that he likes the taste of it and finds it tastier than any of the other options on the app), but which he knows to be prudentially better for him to eat. (Suppose also, here and throughout, that he knows that simple moral vegetarianism is true.) Call any such case in the series *Chain’s End*.

As in Late Stage, morality and prudence conflict in Chain’s End. I take it that *mutatis mutandis*, given simple moral vegetarianism, the following claim about the morality/prudence interaction holds for cases in the series of Meals:

*Morality Always Overrides Prudence In Cases Of Conflict In Meals:* For any case in the series of Meals, if it is prudentially better for Tomaso to eat the relevant meal but not morally permissible to do so, then not eating the meal is (all-things-considered) better than eating it.

To see why, suppose for *reductio* that this alternative claim holds:

*Morality Sometimes Does Not Override Prudence In Cases of Conflict In Meals:* For some case in the series of Meals, if it is prudentially better for Tomaso to eat the relevant meal but not morally

permissible to do so, then eating the meal is *either* (all-things-considered) better than *or* (all-things-considered) equally as good as not eating it.

Suppose also that Tomaso wants to eat the meal in the series of Meals for which this case holds, which entails that he knows that eating the meal is either (all-things-considered) better than or (all-things-considered) equally as good as not eating it. Then practical rationality *permits* him to eat the meal, even though he knows that eating it is speciesist. I take this to be a *reductio* of Morality Sometimes Does Not Override Prudence In Cases of Conflict In Meals. Thus, in Chain's End, *not* eating that meal is (all-things-considered) *better than* eating it.

### 3.2.2. *Chain's Beginning*

Suppose that Tomaso wants to eat a meal in the series which he knows to be morally permissible for him to eat for his reason, and which he knows to be prudentially better for him to eat. Call any such case in the series *Chain's Beginning*.

As in Early Stage, there is no conflict between morality and prudence in Chain's Beginning. I take it that *mutatis mutandis*, given simple moral vegetarianism, the following claim holds for cases in the series of Meals:

*Always Better Be Prudent In Cases Of No Conflict In Meals:* For any case in the series of Meals, if it is prudentially better for Tomaso to eat the relevant meal and morally permissible to do so, then eating the meal is (all-things-considered) better than not eating it.

Thus, in Chain's Beginning, eating that meal is (all-things-considered) *better than* not eating it.

### 3.2.3. *Tomaso's is a case of symmetric vague betterness*

The upshot of these two cases is that, if it is prudentially better for Tomaso to eat a certain meal in the series but *vague* whether morality permits doing so—as in Tomaso's moral

vagueness case—then it is vague whether eating that meal is (all-things-considered) better than not eating it, *and* vague whether not eating it is (all-things-considered) better than eating it. Tomaso’s case is also a case of symmetric vague betterness.

### 3.3. *The Choice Under Symmetric Vague Betterness Question*

Given what I have just argued, the Choice Under Moral Vagueness Question for Francesca’s and Tomaso’s cases can be understood as a special version of a more general question:

#### *Choice Under Symmetric Vague Betterness Question:*

What is it that one rationally ought to do if one knows that it is vague whether  $\varphi$ ing is (all-things-considered) better than not  $\varphi$ ing, and that it is vague whether not  $\varphi$ ing is (all-things-considered) better than  $\varphi$ ing?

## 4. DECISION-MAKING UNDER SYMMETRIC VAGUE BETTERNESS

In Section 3, I showed that Francesca’s and Tomaso’s moral vagueness cases are cases of symmetric vague betterness. In this section, I develop two accounts of decision-making under these cases of symmetric vague betterness. The plan is as follows. In Section 4.1, I show that *any* account of decision-making must answer the following questions in the same way:

#### *Francesca’s Choice Under Symmetric Vague Betterness Question:*

What is it that Francesca rationally ought to do if she knows that it is vague whether aborting is (all-things-considered) better than not aborting, and that it is vague whether not aborting is (all-things-considered) better than aborting?

#### *Tomaso’s Choice Under Symmetric Vague Betterness Question:*

What is it that Tomaso rationally ought to do if he knows that it is vague whether eating a particular meal is (all-things-considered) better than not eating it, and that it is vague whether not eating the meal is (all-things-considered) better than eating it?



In other words, if an account of decision-making holds that Francesca rationally ought (not) to abort, then such an account must hold that Tomaso rationally ought (not) to eat the meal—and vice versa. Why? Because, as I show in Section 4.1, Francesca’s and Tomaso’s cases are structurally identical. In Sections 4.2–3, I argue that vagueness-as-ignorance delivers *one* answer to Francesca’s and Tomaso’s Choice Under Symmetric Vague Betterness Questions. In Sections 4.4–5, I argue that vagueness-as-absence-of-truth delivers *another* answer to those questions. The upshot is that these two influential theories of vagueness yield different accounts of decision-making under symmetric vague betterness.

#### 4.1. *Mistakes, moral mistakes, prudential mistakes*

In developing my account of decision-making under *asymmetric* vague betterness in Chapter 1, I made two assumptions that I also make in what follows. I assume that the relevant rational agents (in this case, Francesca and Tomaso) know what is going on in their choice situations. In other words, I assume that Francesca knows that it is prudentially better to abort, that it is vague whether morality permits her to abort, and so on. *Mutatis mutandis* for Tomaso. Moreover, I assume this definition of what it is to make a mistake in choice (for any choice between  $\varphi$ ing and not  $\varphi$ ing):

*Mistake:*  $\Phi$ ing would be a mistake iff  $\varphi$ ing is all-things-considered worse than not  $\varphi$ ing.<sup>59</sup>  
(*Mutatis mutandis* for not  $\varphi$ ing.<sup>60</sup>)

Given these two assumptions, it is easy to see that, because Francesca knows that she is in the zone of vagueness of the series of Abortions,

---

<sup>59</sup> To be clear, ‘ $\Phi$ ’ is simply the uppercase of ‘ $\varphi$ ’.

<sup>60</sup> That is: Not  $\varphi$ ing would be a mistake iff not  $\varphi$ ing is all-things-considered worse than  $\varphi$ ing.

$A_1, A_2, \dots$ *Zone of vagueness* $\dots, A_{n-1}, A_n$ 


---

SERIES OF ABORTIONS

she knows that it is *vague* whether aborting (henceforth, call this option ‘Aborting’) would be a mistake *and* that it is *vague* whether not aborting (henceforth, ‘Carrying-the-pregnancy-to-term’) would be a mistake. Similarly, it is easy to see that, because Tomaso knows that he is in the zone of vagueness of the series of Meals,

 $M_1, M_2, \dots$ *Zone of vagueness* $\dots, M_{n-1}, M_n$ 


---

SERIES OF MEALS

he knows that it is *vague* whether eating the relevant meal (‘Eating-the-meal’) would be a mistake *and* that it is *vague* whether not eating the meal (‘Refraining-from-eating-it’) would be a mistake. To put it another way, they both know that their choice situations are such that there is vagueness about which of these propositions holds: *Φing is all-things-considered worse than not Φing*, and *Not Φing is all-things-considered worse than Φing*.

However, this is *not* all that Francesca and Tomaso know about their respective choice situations. To begin to see this, let me introduce the two kinds of mistake that are relevant to their choice situations:

*Moral mistake:*  $\Phi$ ing would be a moral mistake iff (1)  $\varphi$ ing is all-things-considered worse than not  $\varphi$ ing, and (2) at least part of what makes it the case that  $\varphi$ ing is all-things-considered worse than not  $\varphi$ ing is that  $\varphi$ ing is *morally* worse than not  $\varphi$ ing. (*Mutatis mutandis* for not  $\varphi$ ing.<sup>61</sup>)

*Prudential mistake:*  $\Phi$ ing would be a prudential mistake iff (1)  $\varphi$ ing is all-things-considered worse than not  $\varphi$ ing, and (2) at least part of what makes it the case that  $\varphi$ ing is all-things-considered worse than not  $\varphi$ ing is that  $\varphi$ ing is *prudentially* worse than not  $\varphi$ ing. (*Mutatis mutandis* for not  $\varphi$ ing.<sup>62</sup>)

Given these definitions, how do we figure out whether a mistake in choice is, for example, a moral mistake? First, we focus our attention both on *all* the considerations that count in favour of  $\varphi$ ing's being good to do and on *all* the considerations that count in favour of not  $\varphi$ ing's being good to do, which may include both moral and non-moral considerations.<sup>63</sup> If, after doing this, we establish that, say,  $\varphi$ ing is (all-things-considered) worse than not  $\varphi$ ing, we then turn our attention to the considerations that count in favour of not  $\varphi$ ing that win out to make it the case that  $\varphi$ ing is (all-things-considered) worse than not  $\varphi$ ing. If these considerations are moral considerations (such that  $\varphi$ ing is morally worse than not  $\varphi$ ing), then  $\varphi$ ing would be a moral mistake. (*Mutatis mutandis* for how to figure out whether a mistake in choice is a prudential mistake.)<sup>64</sup>

One might object that it is not clear why we should accept this definition of moral mistake (and the corresponding definition of prudential mistake), since there is an obvious alternative to it, namely:

---

<sup>61</sup> That is: Not  $\varphi$ ing would be a moral mistake iff not  $\varphi$ ing is all-things-considered worse than  $\varphi$ ing, and at least part of what makes it the case that not  $\varphi$ ing is all-things-considered worse than  $\varphi$ ing is that not  $\varphi$ ing is morally worse than  $\varphi$ ing.

<sup>62</sup> That is: Not  $\varphi$ ing would be a prudential mistake iff not  $\varphi$ ing is all-things-considered worse than  $\varphi$ ing, and at least part of what makes it the case that not  $\varphi$ ing is all-things-considered worse than  $\varphi$ ing is that not  $\varphi$ ing is prudentially worse than  $\varphi$ ing.

<sup>63</sup> Where a moral consideration is any consideration that counts in favour of an action's being morally good to do, morally bad to do, morally required, (merely) morally permissible, or morally impermissible.

<sup>64</sup> In making these points, I draw partially on Elizabeth Harman's [2015: 232; 2016: 368-9] own definition of moral mistake.

*Moral mistake\**:  $\Phi$ ing would be a moral mistake iff  $\varphi$ ing is morally worse than not  $\varphi$ ing. (*Mutatis mutandis* for not  $\varphi$ ing: Not  $\varphi$ ing would be a moral mistake iff not  $\varphi$ ing is morally worse than  $\varphi$ ing.)

Why not define moral mistake in this way rather than in the way that I define it? Because this alternative definition of moral mistake entails that failing to perform *any* morally supererogatory act would be a moral mistake, and this is implausible. To see that the definition entails this, suppose that one is in a choice situation such that  $\varphi$ ing is morally better than not  $\varphi$ ing, but that both  $\varphi$ ing and not  $\varphi$ ing are morally permissible—the features of any case of moral supererogation.<sup>65</sup> Given the alternative definition, failing to  $\varphi$  (i.e., failing to perform an arbitrary morally supererogatory act) would be a moral mistake, since not  $\varphi$ ing is morally worse than  $\varphi$ ing. This contrasts with my definition of moral mistake, which is compatible with all the relevant possibilities: that failing to perform any morally supererogatory act would *not* be a moral mistake, that failing to perform any morally supererogatory act would be a moral mistake (the claim that the alternative definition entails), and that failing to perform *some*, but not any, morally supererogatory act would be a moral mistake (a surprising claim for which Elizabeth Harman [2015, 2016] has recently argued). This is why I define moral mistake (and, correspondingly, prudential mistake) in the way I do.

#### 4.1.1. *What Francesca also knows about her choice situation*

Now consider, to start with her case, the choice situation in which Francesca finds herself. She is deliberating whether to abort—while in the zone of vagueness of the series of Abortions—and so has two options: Aborting and Carrying-the-pregnancy-to-term. Suppose that she chooses Aborting. Then, as we have just seen, it is vague whether she has made a mistake. With the distinction between moral and prudential mistakes in hand, however, we can ask:

---

<sup>65</sup> I say *moral* supererogation because there might be supererogation across normative domains. For an argument that there is, see McElwee [2017].

*Aborting Question:*

Why is it that, by Aborting, it is vague whether Francesca has made a mistake? Is it because (Answer A) it is *both* vague whether she has made a moral mistake *and* vague whether she has made a prudential mistake?

(Answer B) she has (determinately) *not* made a moral mistake, but it is *vague* whether she has made a prudential mistake?

(Answer C) she has (determinately) *not* made a prudential mistake, but it is *vague* whether she has made a moral mistake?

The right answer to the Aborting Question is C. To see this, consider first why Francesca has not made a *prudential* mistake: simply because Aborting is prudentially better—and hence *not* prudentially worse—than Carrying-the-pregnancy-to-term.

Consider now why it is vague whether Francesca has made a *moral* mistake: because it is vague whether Aborting is all-things-considered worse than Carrying-the-pregnancy-to-term—as argued in Section 3—and *vague whether Aborting is morally worse than Carrying-the-pregnancy-to-term*. Why is the latter claim (in italics) true? Notice first that:

- (1) it is vague whether Aborting is morally permissible for her (as we saw in Section 2 through the relevant moral sorites paradox), and
- (2) Carrying-the-pregnancy-to-term is (determinately) morally permissible for her. (After all, she would be doing nothing wrong, morally, if she had the child.)

Suppose now, for the sake of argument, that there is *no* vagueness about whether Aborting is morally permissible for Francesca. If Aborting is *not* morally permissible for her, then—given that Carrying-the-pregnancy-to-term is morally permissible for her—Carrying-the-pregnancy-to-term would have greater moral value than Aborting, in which case *Aborting would be morally worse than Carrying-the-pregnancy-to-term*. On the other hand, if Aborting *is* morally permissible for her, then neither Aborting nor Carrying-the-pregnancy-to-term would have greater moral

value than the other, in which case *neither option would be morally worse than the other*. Why? Because Aborting and Carrying-the-pregnancy-to-term would both be morally permissible *and* neither of these options is, as a matter of fact, morally supererogatory. Therefore, because there *is* vagueness about whether Aborting is morally permissible for Francesca, and Carrying-the-pregnancy-to-term is (determinately) morally permissible for her, it is indeed vague whether Aborting is morally worse than Carrying-the-pregnancy-to-term.

The upshot of this is that, since Francesca knows what is going on in her choice situation, she knows not merely that if she chooses Aborting, it is vague whether she would be making a mistake. She *also* knows that if she chooses Aborting, she would (determinately) *not* be making a *prudential* mistake, and that it is *vague* whether she would be making a *moral* mistake.

Let us now consider the other possibility: that Francesca chooses Carrying-the-pregnancy-to-term. If she does this, then, as we saw, it is also vague whether she has made a mistake, and we can similarly ask:

*Carrying-the-pregnancy-to-term Question:*

Why is it that, by Carrying-the-pregnancy-to-term, it is vague whether Francesca has made a mistake? Is it because

(**Answer A**) it is *both* vague whether she has made a moral mistake *and* vague whether she has made a prudential mistake?

(**Answer B**) she has (determinately) *not* made a moral mistake, but it is *vague* whether she has made a prudential mistake?

(**Answer C**) she has (determinately) *not* made a prudential mistake, but it is *vague* whether she has made a moral mistake?

The right answer to the Carrying-the-pregnancy-to-term Question is B. To see this, consider first why Francesca has not made a moral mistake: because *Carrying-the-pregnancy-to-term is (determinately) not morally worse than Aborting*. Why is this? Recall that it is vague whether Aborting

is morally permissible for Francesca, and that Carrying-the-pregnancy-to-term is (determinately) morally permissible for her. This implies—as we have just seen in discussing the Aborting Question—that if there were no vagueness about whether Aborting is morally permissible for Francesca, then either Aborting would be morally worse than Carrying-the-pregnancy-to-term or neither option would be morally worse than the other. Thus, in no scenario is Carrying-the-pregnancy-to-term morally worse than Aborting.

Consider now why it is vague whether Francesca has made a prudential mistake: because it is vague whether Carrying-the-pregnancy-to-term is all-things-considered worse than Aborting—as argued in Section 3—and *Carrying-the-pregnancy-to-term is (determinately) prudentially worse than Aborting*. The reason why the latter claim holds is straightforward: as we have just seen in discussing the Aborting Question, Aborting is prudentially better than Carrying-the-pregnancy-to-term, which implies that Carrying-the-pregnancy-to-term is prudentially worse.

As in the Aborting Question scenario, the upshot of this is similarly that, since Francesca knows what is going on in her choice situation, she knows not merely that if she chooses Carrying-the-pregnancy-to-term, it is vague whether she would be making a mistake. She *also* knows that if she chooses Carrying-the-pregnancy-to-term, she would (determinately) *not* be making a *moral* mistake, and that it is *vague* whether she would be making a *prudential* mistake.

To make the answers to the Aborting and Carrying-the-pregnancy-to-term Questions clearer, I summarise them below in what I call a ‘table of mistakes’ (which also includes, to enhance clarity, the other cases in the series of Abortions), where ‘ $\nabla$ ’ stands for ‘vague whether’, ‘ $\Delta$ ’ for ‘determinately’, and ‘ $\neg$ ’ for ‘not’:

	$A_1, A_2, \dots$	<i>Zone of vagueness</i>	$\dots, A_{n-1}, A_n$
<i>Aborting</i>	$\Delta\neg(\text{A moral mistake})$ & $\Delta\neg(\text{A prudential mistake})$	$\nabla(\text{A moral mistake})$ & $\Delta\neg(\text{A prudential mistake})$	$\Delta(\text{A moral mistake})$ & $\Delta\neg(\text{A prudential mistake})$
<i>Carrying-the-pregnancy-to-term</i>	$\Delta\neg(\text{A moral mistake})$ & $\Delta(\text{A prudential mistake})$	$\Delta\neg(\text{A moral mistake})$ & $\nabla(\text{A prudential mistake})$	$\Delta\neg(\text{A moral mistake})$ & $\Delta\neg(\text{A prudential mistake})$

TABLE OF MISTAKES I: Moral and prudential mistakes in the series of Abortions

#### 4.1.2. *What Tomaso also knows about his choice situation*

Let us now turn to the choice situation in which Tomaso finds himself. As I show in what follows, his situation is structurally identical to Francesca's. Tomaso is deliberating whether to eat a meal—a meal which is in the zone of vagueness of the series of Meals—and so has two options: Eating-the-meal and Refraining-from-eating-it. Suppose that he chooses Eating-the-meal. Then, as we saw, it is vague whether he has made a mistake. Given the distinction between moral and prudential mistakes, we can ask:

*Eating-the-meal Question:*

Why is it that, by Eating-the-meal, it is vague whether Tomaso has made a mistake? Is it because

(**Answer A**) it is *both* vague whether he has made a moral mistake *and* vague whether he has made a prudential mistake?

(**Answer B**) he has (determinately) *not* made a moral mistake, but it is *vague* whether he has made a prudential mistake?

(**Answer C**) he has (determinately) *not* made a prudential mistake, but it is *vague* whether he has made a moral mistake?

The right answer to the Eating-the-meal Question is C. Why? Consider first why Tomaso has not made a prudential mistake: because Eating-the-meal is prudentially better—and hence *not* prudentially worse—than Refraining-from-eating-it.



Consider now why it is vague whether Tomaso has made a moral mistake: because it is vague whether Eating-the-meal is all-things-considered worse than Refraining-from-eating-it—as argued in Section 3—and *it is vague whether Eating-the-meal is morally worse than Refraining-from-eating-it*. Why is the latter claim true? Notice first that:

- (1) it is vague whether Eating-the-meal is morally permissible for him (as we saw in Section 2 through the relevant moral sorites paradox), and
- (2) Refraining-from-eating-it is (determinately) morally permissible for him. (After all, he would be doing nothing wrong, morally, if he refrained from eating that meal.)

Suppose now, for the sake of argument, that there is *no* vagueness about whether Eating-the-meal is morally permissible for Tomaso. If Eating-the-meal is *not* morally permissible for him, then—given that Refraining-from-eating-it is morally permissible for him—Refraining-from-eating-it would have greater moral value than Eating-the-meal, in which case *Eating-the-meal would be morally worse than Refraining-from-eating-it*. On the other hand, if Eating-the-meal *is* morally permissible for him, then neither Eating-the-meal nor Refraining-from-eating-it would have greater moral value than the other, in which case *neither option would be morally worse than the other*. Why? Because Eating-the-meal and Refraining-from-eating-it would both be morally permissible *and* neither of these options is, as a matter of fact, morally supererogatory. Therefore, because there *is* vagueness about whether Eating-the-meal is morally permissible for Tomaso, and Refraining-from-eating-it is (determinately) morally permissible for her, it is indeed vague whether Eating-the-meal is morally worse than Refraining-from-eating-it.

Let us now suppose that Tomaso chooses Refraining-from-eating-it. If he does this, then it is also vague whether he has made a mistake, and we can similarly ask:

*Refraining-from-eating-it Question:*

Why is it that, by Refraining-from-eating-it, it is vague whether Tomaso has made a mistake?

Is it because

(**Answer A**) it is *both* vague whether he has made a moral mistake *and* vague whether he has made a prudential mistake?

(**Answer B**) he has (determinately) *not* made a moral mistake, but it is *vague* whether he has made a prudential mistake?

(**Answer C**) he has (determinately) *not* made a prudential mistake, but it is *vague* whether he has made a moral mistake?

The right answer to the Refraining-from-eating-it Question is B. To see this, consider first why Tomaso has not made a moral mistake: because *Refraining-from-eating-it is (determinately) not morally worse than Aborting*. Why is this? Recall that it is vague whether Eating-the-meal is morally permissible for Tomaso, and that Refraining-from-eating-it is (determinately) morally permissible for him. This implies—as we have just seen in discussing the Eating-the-meal Question—that if there were no vagueness about whether Eating-the-meal is morally permissible for Tomaso, then either Eating-the-meal would be morally worse than Refraining-from-eating-it or neither option would be morally worse than the other. So, in no scenario is Refraining-from-eating-it morally worse than Eating-the-meal.

Consider now why it is vague whether Tomaso has made a prudential mistake: because it is vague whether Refraining-from-eating-it is all-things-considered worse than Eating-the-meal—as argued in Section 2—and *Refraining-from-eating-it is (determinately) prudentially worse than Eating-the-meal*. Why is the latter claim true? Because, as we have just seen in discussing the Eating-the-meal Question, Eating-the-meal is prudentially better, and so not prudentially worse, than Refraining-from-eating-it.

As should by now be clear, these arguments are structurally identical to the ones in §4.1.1. Their upshot is also structurally identical: given that Tomaso knows what is going on in his choice situation, he knows not simply that, no matter what he does (whether he chooses

Eating-the-meal or Refraining-from-eating-it), it is vague whether he would be making a mistake. He *also* knows that if he chooses Eating-the-meal, he would (determinately) *not* be making a *prudential* mistake and that it is *vague* whether he would be making a *moral* mistake; and he knows that if he chooses Refraining-from-eating-it, he would (determinately) *not* be making a *moral* mistake and that it is *vague* whether he would be making a *prudential* mistake. I summarise these answers to the Eating-the-meal and Refraining-from-eating-it Questions in the table of mistakes for the series of Meals:

	$M_1, M_2, \dots$	<i>Zone of vagueness</i>	$\dots, M_{n-1}, M_n$
<i>Eating-the-meal</i>	$\Delta \neg(\text{A moral mistake})$ & $\Delta \neg(\text{A prudential mistake})$	$\nabla(\text{A moral mistake})$ & $\Delta \neg(\text{A prudential mistake})$	$\Delta(\text{A moral mistake})$ & $\Delta \neg(\text{A prudential mistake})$
<i>Refraining-from-eating-it</i>	$\Delta \neg(\text{A moral mistake})$ & $\Delta(\text{A prudential mistake})$	$\Delta \neg(\text{A moral mistake})$ & $\nabla(\text{A prudential mistake})$	$\Delta \neg(\text{A moral mistake})$ & $\Delta \neg(\text{A prudential mistake})$

TABLE OF MISTAKES II: Moral and prudential mistakes in the series of Meals

#### 4.1.3. *The upshot so far*

The conclusion to draw from the preceding discussion is that, since Francesca's and Tomaso's cases of symmetric vague betterness are structurally identical, any account of decision-making under such cases must deliver a similar answer to both Francesca's and Tomaso's Choice Under Symmetric Vague Betterness Questions, which can now be more accurately understood as asking:

*Francesca's Choice Under Symmetric Vague Betterness Question:*

What is it that Francesca rationally ought to do if she knows that (1) Aborting would (determinately) not be a prudential mistake but that it is vague whether it would be a moral mistake, and that (2) Carrying-the-pregnancy-to-term would (determinately) not be a moral mistake but that it is vague whether it would be a prudential mistake?

*Tomaso's Choice Under Symmetric Vague Betterness Question:*

What is it that Tomaso rationally ought to do if he knows that (1) Eating-the-meal would (determinately) not be a prudential mistake but that it is vague whether it would be a moral mistake, and that (2) Refraining-from-eating-it would (determinately) not be a moral mistake but that it is vague whether it would be a prudential mistake?

#### 4.2. *Vagueness-as-ignorance*

According to vagueness-as-ignorance, vagueness is a kind of ignorance. To flesh out vagueness-as-ignorance using a straightforward example from Chapter 1, suppose that Brad Pitt, who is 1800 mm tall, is a borderline case of the predicate 'is tall' (for a male adult). On vagueness-as-ignorance, there is a fact of the matter about whether Brad Pitt is tall, but it is unknowable—or, as I prefer to put it, there is an *unknowable fact of the matter* about whether he is tall. Accordingly, on this view, Brad Pitt is either tall or not tall and the proposition *Brad Pitt is tall* is either true or false (i.e., both the law of excluded middle<sup>66</sup> and the law of bivalence<sup>67</sup>, respectively, hold). To put it another way, 'is tall' draws, contrary to appearances, a sharp boundary between the men who are tall and those who are not tall, and so either Brad Pitt falls on the 'tall' side of that boundary (in which case *Brad Pitt is tall* is true) or he falls on the 'not tall' side of it (in which case *Brad Pitt is tall* is false). However, the location of that boundary is unknowable.

But if there is a fact of the matter about whether Brad Pitt is tall, why is it that we are not able to know it? What explains, in other words, that the location of the sharp boundaries drawn by vague predicates—of which 'is tall' is merely an example—be unknowable? The standard explanation is Timothy Williamson's [1994: 217-47].

---

<sup>66</sup> *Law of excluded middle*: For any proposition  $p$ , either  $p$  or not- $p$ .

<sup>67</sup> *Law of bivalence*: For any proposition  $p$ , either  $p$  is true or  $p$  is false.

Williamson's explanation is a safety-based explanation, appealing to what has come to be referred to as the safety condition for knowledge, according to which, in Williamson's own words: 'If one knows, one could not easily have been wrong in a similar case' [Williamson 2000: 147]. A bit more precisely, the condition states that

*Safety*: If one knows that  $p$ , then there is no close possible world in which one has a sufficiently similar belief (let us say, a belief that  $p^*$ ) and it is false that  $p^*$ ,

where a close possible world can be distinguished from a far-off possible world by way of two simple examples: the world in which I move the laptop where I am writing these words five centimetres to my right, and the world in which Lisbon and New York City have swapped places with each other. The former is a close possible world in that it is close to the actual world, where the closeness is determined by the similarity between the two worlds—in this case, by how much the actual world needs to change for it to be like the possible world in which I move the laptop five centimetres to my right.<sup>68</sup> Clearly, it does not need to change much—hence why that possible world is close. Conversely, the world in which Lisbon and New York City have swapped places with each other is a far-off possible world in that much needs to change in the actual world for it to be like this possible world.

To understand Williamson's safety-based explanation, let us assume that, as a matter of fact, the boundary of 'is tall' is located at 1800 mm and that we believe this. Accordingly, since we know Brad Pitt's height (1800 mm), we believe *Brad Pitt is tall* and this belief is true. However, given Safety, if there are close possible worlds in which we have a sufficiently similar belief (a belief that *Brad Pitt is tall\**) and it is false that *Brad Pitt is tall\**, then our actual belief (i.e., that *Brad Pitt is tall*) does not constitute knowledge. Are there such worlds? According to Williamson, there are. Had the overall use of 'is tall' been very slightly different, its boundary

---

<sup>68</sup> For a defence of modal closeness in terms of similarity, see Lewis [1987].

could easily have been located at, say, 1801 mm and we could easily have still believed that it is located at 1800 mm. Its boundary could easily have been located at 1801 mm because, Williamson holds, the boundaries of vague predicates are not just sharp but also *unstable*: very slight shifts in their use could bring about very slight shifts in their boundaries.<sup>69</sup> And we could easily have still believed that the boundary of ‘is tall’ is located at 1800 mm because, Williamson holds, the capacity of humans to discriminate between shifts that would locate the boundary at 1800 mm and shifts that would locate it at 1801 mm is limited.

Consider then the close possible world *W* in which the boundary of ‘is tall’ is located at 1801 mm, and assume that all the other relevant facts are held constant in *W*: that Brad Pitt exists, that he is 1800 mm tall, that we know his height, that we believe that the boundary of ‘is tall’ is located at 1800 mm, and so on. In *W*, we have a very similar belief to *Brad Pitt is tall*—namely, *Brad Pitt is tall\**—which we express using the same words as in the actual world: ‘Brad Pitt is tall’. However, because ‘is tall’ (for a male adult), in *W*, refers to the property of being at least 1801 mm tall (and not to the property of being at least 1800 mm tall, as in the actual world), our belief that *Brad Pitt is tall\** is false. By Safety, it follows that we do not know that *Brad Pitt is tall*—and in practice, due to our limited capacities, we cannot know it. After all, even though our belief that *Brad Pitt is tall* is true, it is true as a matter of luck: we could *easily* have been wrong in believing it.<sup>70</sup> In a nutshell, the belief is true but not safely true. This is Williamson’s version of vagueness-as-ignorance.

Though Williamson’s is the standard version of vagueness-as-ignorance, it is problematic when applied to moral vagueness. Notice that, given his version of the view, the boundary of ‘is morally permissible’ (like that of any other vague predicate) is unstable: very slight shifts in its pattern of *use* could lead to very slight shifts in the location of its boundary. Recently, Miriam Schoenfield [2016] argued that this feature of Williamson’s version of the

---

<sup>69</sup> This instability of vague predicates on Williamson’s account is sometimes referred to as their semantic plasticity.

<sup>70</sup> For more on ‘easy’ possibilities, see Sainsbury [1997].

view has implausible implications for moral deliberation. Consider its implications, for example, for Francesca’s moral vagueness case. Williamson’s version of the view implies that, as Schoenfield puts it, Francesca ‘could (in principle, though it would be extremely difficult!) resolve her deliberation about whether aborting her fetus is permissible by learning enough about her community’s linguistic usage’ [Schoenfield 2016: 277].<sup>71</sup> For this reason, and because the other versions of vagueness-as-ignorance in the literature also seem to face problems when applied to moral vagueness<sup>72</sup>, I do not assume any particular version of the view in what follows. Instead, in what follows, I assume merely the basic tenets of the view (as spelt out in the first paragraph of this section, Section 4.2), with the aim of figuring out, conditional on the view, what account of decision-making it delivers under our cases of symmetric vague betterness.

#### 4.3. *The Play-it-Safe account of decision-making under symmetric vague betterness*

Let us assume then vagueness-as-ignorance. To make clear the implications of it for our cases of symmetric vague betterness, consider the series of Abortions, which starts with cases in which Aborting is (all-things-considered) better than Carrying-the-pregnancy-to-term and ends with cases in which Carrying-the-pregnancy-to-term is (all-things-considered) better than Aborting. Given vagueness-as-ignorance, there is an abortion case,  $A_i$ , in the series of Abortions such that  $A_i$  is the *last* case in which Aborting is (all-things-considered) better than Carrying-the-pregnancy-to-term, and, correspondingly, there is a case in the series—namely,

---

<sup>71</sup> This objection to Williamson’s version of vagueness-as-ignorance is part of an argument of Schoenfield’s to the conclusion that if a robust form of moral realism is true (and there is moral vagueness), then moral vagueness is—at least in part—a metaphysical phenomenon, as opposed to a purely semantic or epistemic phenomenon. (See also Manley [MS], who appears to have independently identified the same problem as Schoenfield for Williamson’s version of the view.)

<sup>72</sup> Schoenfield [2016: 278-281] argues persuasively that, when applied to moral vagueness, Kearns and Magidor’s [2008] version of the view has similar implications for moral deliberation, and that Roy Sorensen’s [2001] version would require that moral truths be contingently true. (Interestingly, Patrick Greenough [2008: 229] thinks that Sorensen’s version, despite how it is discussed in the literature and by Sorensen himself, is not a version of vagueness-as-ignorance at all.) For more general objections to Sorensen’s version of the view, see Williamson [2007] and Jago [2013].

the next one,  $A_i + 1$ —such that  $A_i + 1$  is the *first* case in which Carrying-the-pregnancy-to-term is (all-things-considered) better than Aborting. In other words, there is a sharp boundary between the cases in which Aborting is better and those in which Carrying-the-pregnancy-to-term is better, since the predicate ‘is morally permissible’—the source of the vagueness in ‘(all-things-considered) better than’—draws a sharp boundary. *Mutatis mutandis* for the series of Meals. Furthermore, given vagueness-as-ignorance, it is unknowable where  $A_i$  and  $A_i + 1$  are located—hence the zone of vagueness in both series. This implies that, if one acts one way or the other in the zone of vagueness of any of the series, there is an unknowable fact of the matter about whether one has made a mistake in choice.

With all these implications of vagueness-as-ignorance in mind, consider (to start with her case) Francesca’s Choice Under Symmetric Vague Betterness Question:

What is it that Francesca rationally ought to do if she knows that (1) Aborting would (determinately) not be a prudential mistake but that it is vague whether it would be a moral mistake, and that (2) Carrying-the-pregnancy-to-term would (determinately) not be a moral mistake but that it is vague whether it would be a prudential mistake?

(Recall that Francesca knows (1) and (2).) To begin with, notice that *no matter what she does*, whether she chooses Aborting or Carrying-the-pregnancy-to-term, it could *easily* turn out that, by bad luck, she is making *some* mistake. Before explaining why this is the case, let us assume that the boundary of moral permissibility is located at  $A_i$ , such that  $A_i$  is the last case (or point) in the series in which it is morally permissible for Francesca to abort and  $A_i + 1$  is the first case (point) in the series in which it is morally impermissible for her to do so. Thus, if Francesca is at a point in the series before  $A_i + 1$  and makes up her mind to carry the pregnancy to term, she is making a prudential mistake; if she is at a point in the series after  $A_i$  and aborts, she is making a moral mistake. Here is then why—in the zone of vagueness—it could easily turn out that, by bad luck, Francesca is making some mistake regardless of whether she chooses



Aborting or Carrying-the-pregnancy-to-term. On the one hand, if she chooses Aborting, it could easily turn out that, by bad luck, she is making a moral mistake. Why? Because even if it turns out that she did not make a moral mistake by Aborting (i.e., she aborted at a point before  $A_i + 1$ —before it became morally impermissible to abort), it could easily have turned out that—there is a close possible world in which—she did make a moral mistake by Aborting, namely, had she aborted a little later than she actually did (viz., after  $A_i$ —after it became morally impermissible to abort). On the other hand, if she chooses Carrying-the-pregnancy-to-term, it could easily turn out that, by bad luck, she is making a prudential mistake. Why? Because even if it turns out that she did not make a prudential mistake by Carrying-the-pregnancy-to-term (i.e., she made up her mind to carry the pregnancy to term at a point after  $A_i$ ), it could easily have turned out that—there is a close possible world in which—she did make a prudential mistake by Carrying-the-pregnancy-to-term, namely, had she made up her mind to carry the pregnancy to term a little earlier than she actually did (viz., before  $A_i + 1$ ).

Whatever Francesca does, then, it is a matter of *luck* whether she is making a mistake in choice. This amounts to there being a *high risk* that she is making some mistake whatever she does, at the very least if risk is understood modally. Let me say some more about risk understood modally in the next three paragraphs, before returning to the question of what Francesca rationally ought to do.

To talk of risk is to say, at minimum, that an event—in this case, making some mistake in choice—is a possible but unwanted outcome of some action (in this case, of the action of Aborting at a certain point in the series, and of the action of Carrying-the-pregnancy-to-term at a certain point in the series). Call this a *risk event* (following Pritchard [2015: 437]). In one way of understanding risk, a risk event is to be understood in terms of probability. Take, for example, the relatively low risk of dying in a plane crash. Why is flying on a plane a relatively low-risk activity? In the probabilistic way of understanding risk, because the probability of the risk event occurring—where the risk event, in this case, is dying in a plane crash—is generally

low, and it is not significantly high when compared to the risk of dying using other means of transportation [Pritchard 2015: 440].

In another way of understanding risk, a risk event is to be understood in terms of modal closeness [Pritchard 2015; see also Williamson 2009]. To see that a modal understanding of risk can come apart from a probabilistic one, consider the following (deliberately extreme) case. Suppose that, unbeknownst to everyone, a terrorist has hidden a bomb in Stockholm that, if detonated, will kill and severely injure many people, but that there is something bizarre about this act of terrorism: if the lottery ticket that the terrorist has just bought is a jackpot *loser*, the bomb will deactivate by itself; otherwise, if the ticket is a jackpot *winner*, the bomb will detonate.<sup>73</sup> The odds of any of those lottery tickets being a jackpot winner are (say) 1 in 10,000,000. So, the risk event—the bomb detonating—has a low probability of occurring.

However, as Duncan Pritchard [2015: 443] points out, that is consistent with the risk event being an event that could *easily* occur. To put it another way, it is consistent with there being a *close* possible world in which the ticket is a jackpot winner and the bomb detonates. Why? Because the actual world need not change much for it to be like this possible world. For that to happen, it is just a matter of a few balls coming out of a machine to be slightly differently numbered. In a modal way of understanding risk, then, a risk event that occurs in a close possible world is a *high*-risk event (and, correspondingly, a risk event that occurs in a far-off possible world is a low-risk event).

This is what I mean when I say—to return now to the question of what Francesca rationally ought to do—that there is a high risk that she is making some mistake regardless of whether she chooses Aborting or Carrying-the-pregnancy-to-term: there is a high risk (at the very least) in a modal understanding of risk. Since Francesca knows that she is in the zone of vagueness, she knows that there is a high risk (modally speaking) of making a moral mistake

---

<sup>73</sup> Adapted from a case by Pritchard [2015: 441].

by Aborting and that there is a high risk (modally speaking) of making a prudential mistake by Carrying-the-pregnancy-to-term. Thus, Francesca ought to believe that:

(i) *There is a high risk (modally speaking) of making a moral mistake by Aborting.*

Now, here is a platitude about mistakes in choice: some mistakes are more serious than others. In Francesca's case, I take it that making *that* moral mistake would be substantially more serious than making *that* prudential mistake. After all, performing an act that deprives a sentient (human) foetus of a life worth living would be a substantially more serious mistake than performing an act that compromises one's career plans. (Recall that making that moral mistake would amount to killing a sentient foetus because, given the simple liberal view, the first case in the series of Abortions in which it is morally impermissible for Francesca to abort is also the first case in the series in which the foetus is sentient. So, if Francesca aborted for her reason at that point in the series, she would be doing something very morally wrong.) Since Francesca knows the choice situation that she is in, she ought to believe that:

(ii) *Making that moral mistake would be substantially more serious than making that prudential mistake.*

Moreover, she also ought to believe that:

(iii) *Carrying-the-pregnancy-to-term would not be a moral mistake.*

By Aborting, then, she would be running an unjustified risk. By *not* Aborting, however, she would be certain to not be making a moral mistake. So, she rationally ought to play it safe and choose Carrying-the-pregnancy-to-term.

*Mutatis mutandis*, the same can be said in response to Tomaso's Choice Under Symmetric Vague Betterness Question:

What is it that Tomaso rationally ought to do if he knows that (1) Eating-the-meal would (determinately) not be a prudential mistake but that it is vague whether it would be a moral mistake, and that (2) Refraining-from-eating-it would (determinately) not be a moral mistake but that it is vague whether it would be a prudential mistake?

Given what he knows, Tomaso ought to believe that:

- (i) *There is a high risk (modally speaking) of making a moral mistake by Eating-the-meal.*
- (ii) *Making that moral mistake (viz., performing an act that results from the prior infliction of pain upon a sentient non-human being) would be substantially more serious than making that prudential mistake (viz., performing an act that deprives oneself of gustatory pleasure, or of greater gustatory pleasure).*
- (iii) *Refraining-from-eating-it would not be a moral mistake.*

By Eating-the-meal, then, Tomaso would be running an unjustified risk. By *not* Eating-the-meal, however, he would be certain to not be making a moral mistake. So, like Francesca, he rationally ought to play it safe, which in his case amounts to choose Refraining-from-eating-it.

If my argument is sound, then, conditional on vagueness-as-ignorance, the right account of decision-making under *these* cases of symmetric vague betterness is that there is *no* vagueness about what one rationally ought to do: one (determinately) ought to play it safe and not  $\varphi$ . I call this the *Play-it-Safe account* of decision-making under symmetric vague betterness.

At this point, one might object that it is not clear what the Play-it-Safe account implies in some choice situations. Suppose that one knows that one is in a choice situation in which—like Francesca’s and Tomaso’s moral vagueness cases—it is (determinately) prudentially better to  $\varphi$  and vague whether it is morally permissible to  $\varphi$ , but in which—unlike Francesca’s and Tomaso’s cases—the relevant prudential considerations are weighty and the relevant moral considerations are flimsy. Does the Play-it-Safe account imply that one rationally ought to not

$\varphi$  in this choice situation? No. The Play-it-Safe account is compatible with the view that—because it is a choice situation in which the prudential considerations are weightier than the moral considerations—this choice situation is such that, regardless of whether  $\varphi$ ing is morally permissible or not,  $\varphi$ ing is (determinately) all-things-considered *better than* not  $\varphi$ ing. In other words, the Play-it-Safe account is compatible with that choice situation being one in which there is *no* vagueness in ‘(all-things-considered) better than’, and so a situation for which the account does not hold (because the account is an account of decision-making only under cases of vagueness in ‘(all-things-considered) better than’).

Only if combined with the following claims would the Play-it-Safe account imply that one rationally ought to not  $\varphi$  in that choice situation:

*Morality Always Overrides Prudence In Cases Of Conflict:* For any choice situation in which only morality and prudence matter, if it is prudentially better to  $\varphi$  but not morally permissible to  $\varphi$ , then not  $\varphi$ ing is (all-things-considered) better than  $\varphi$ ing.

*Always Better Be Prudent In Cases Of No Conflict:* For any choice situation in which only morality and prudence matter, if it is prudentially better to  $\varphi$  and morally permissible to  $\varphi$ , then  $\varphi$ ing is (all-things-considered) better than not  $\varphi$ ing.

But these claims are not being assumed. What is being assumed—because the simple liberal view is being assumed—is that morality overrides prudence whenever there is a conflict between them *in cases in the series of Abortions*, and that it is better to be prudent whenever there is no conflict between morality and prudence *in cases in the series of Abortions*. *Mutatis mutandis* for cases in the series of Meals (because simple moral vegetarianism is being assumed). This is compatible with there being some choice situations in which morality does not override prudence in cases of conflict, such as, presumably, the choice situation that we considered in

the last paragraph. More generally, the Play-it-Safe account holds merely for choice situations, in which only morality and prudence matter, that have the following features:

- (1) it is (determinately) prudentially better to  $\varphi$ ;
- (2) it is vague whether it is morally permissible to  $\varphi$ ;
- (3) not  $\varphi$ ing is (all-things-considered) better than  $\varphi$ ing iff it is not morally permissible to  $\varphi$ ; and
- (4)  $\varphi$ ing is (all-things-considered) better than not  $\varphi$ ing iff it is morally permissible to  $\varphi$ .

If the choice situation that we considered in the last paragraph is indeed one in which  $\varphi$ ing is (determinately) all-things-considered better than not  $\varphi$ ing, then it is a situation that satisfies (1) and (2) but neither (3) nor (4) (because otherwise there would be vagueness in ‘(all-things-considered) better than’), and so a situation for which Play-it-Safe account does not hold.

#### 4.4. *Vagueness-as-absence-of-truth*

According to vagueness-as-absence-of-truth, of which supervaluationism is a version, vagueness is not a kind of ignorance. Instead, vagueness is a kind of *indeterminacy* that results in truth-value gaps. On this view, there is *no* fact of the matter about whether Brad Pitt is tall. More specifically, the proposition *Brad Pitt is tall* is neither true nor false (i.e., the law of bivalence fails), and so ‘is tall’ does not draw a sharp boundary between the men for whom it is true that they are tall and those for whom it is false that they are tall. Like vagueness-as-ignorance, this view also holds that we cannot know that *Brad Pitt is tall*, but for a very different and much more straightforward reason: because knowledge implies truth and *Brad Pitt is tall* is not true.

But what is it to say that vagueness is a kind of indeterminacy that results in truth-value gaps? David Taylor [2018] has developed a minimal theory of indeterminacy on which I want to draw here to answer this question and thereby make vagueness-as-absence-of-truth clearer.

(The theory is ‘minimal’ because it is intended to be compatible with any substantive theory of the nature or logic of indeterminacy.<sup>74</sup>) In what follows, I introduce Taylor’s minimal theory (§4.4.1), and then use it to clarify certain features of vagueness-as-absence-of-truth (§4.4.2).

#### 4.4.1. *What is indeterminacy?*

Plausibly, to say ‘it is indeterminate whether  $p$ ’ is equivalent to saying ‘there is no fact of the matter about whether  $p$ ’. According to Taylor, this equivalence suggests that there is a reasoning that is standard in arguments that appeal to indeterminacy: if anything (metaphysically) determines whether  $p$ , it is facts of a certain type (call them *facts of type B*, where ‘B’ stands for ‘Base’); but these facts do not determine that  $p$  and they do not determine that not- $p$ ; so, it is indeterminate whether  $p$ . The ‘standard reasoning’, as Taylor calls it, applies across the wide range of topics in which indeterminacy arguments have been proposed, of which vagueness is merely one example:

*Translation:* If anything determines whether ‘gavagai’ translates as ‘rabbit’, it is the facts of type B, in this case facts about linguistic behaviour; but these facts do not determine whether ‘gavagai’ translates as ‘rabbit’ or not; so, it is indeterminate whether ‘gavagai’ translates as ‘rabbit’—or so Quine [1960] argued.

*Personal Identity:* If anything determines whether Alpha is the same person as Beta, it is facts of type B, in this case facts about psychological and physical continuity; but these facts do not determine whether Alpha is the same person as Beta or not (in brain fission cases); so, it is indeterminate whether Alpha is the same person as Beta—or so Parfit [1984] argued.

*Vagueness:* If anything determines whether (to use our example) Brad Pitt is tall, it is facts of type B, in this case facts about the height of Brad Pitt; but these facts do not determine whether

---

<sup>74</sup> In this way, Taylor’s project has obvious similarities with the well-established project in the vagueness literature of providing a minimal theory of *vagueness*. For an attempt to provide such a theory, see Greenough [2003].

Brad Pitt is tall or not; so, it is indeterminate whether Brad Pitt is tall—or so those who think that vagueness is a kind of indeterminacy have argued.

Given the standard reasoning, cases of indeterminacy can plausibly be conceived as involving a *determination failure* by a set of facts—there are facts that are supposed (whose ‘job’ it is) to determine whether  $p$ , but which fail to determine that  $p$  and fail to determine that not- $p$ . These observations lead to a preliminary version of the minimal theory of indeterminacy:

*Preliminary MTI*: It is indeterminate whether  $p$  iff there is a type of fact B such that (i) B is a determination base for whether  $p$  and (ii) the B-facts fail to determine whether  $p$  [Taylor 2018: 18],

where a type of fact B is a *determination base* for whether  $p$  iff if anything determines whether  $p$ , it is facts of type B; and where the *B-facts* are the facts of type B that actually obtain (if any do).

For methodological reasons that do not matter for our purposes here (viz., that Preliminary MTI does not meet one of the desiderata for a minimal theory), to arrive at the final version of the minimal theory of indeterminacy, Preliminary MTI should be translated into metalinguistic terms. To that end, Taylor proposes one metalinguistic surrogate for the notion of a determination base and another for the notion of a determination failure. As a surrogate for determination base, he proposes a *sentential* determination base: a type of fact B is a sentential determination base for the sentence ‘ $p$ ’ (given the meaning of ‘ $p$ ’) iff if anything determines whether ‘ $p$ ’ is true, it is facts about the meaning of ‘ $p$ ’ (the M-facts) together with the B-facts [Taylor 2018: 19-20]. As for determination failure (according to which certain facts fail to determine that  $p$  and fail to determine that not- $p$ ), he proposes the following surrogate: certain facts, together with facts about the meaning of ‘ $p$ ’, fail to determine that ‘ $p$ ’ is true and fail to determine that ‘ $p$ ’ is false [Taylor 2018: 20]. With these two surrogates in hand, the final version of the minimal theory can be stated as follows:



*Minimal theory of indeterminacy:* It is indeterminate whether  $p$  iff there is a type of fact B such that (i) B is a sentential determination base for ' $p$ ' (given the meaning of ' $p$ '); and (ii) the B-facts and the facts about the meaning of ' $p$ ' (the M-facts) fail to determine that ' $p$ ' is true, and the B-facts and the M-facts fail to determine that ' $p$ ' is false.

To illustrate this, consider the sentence 'Brad Pitt is tall'. If anything determines whether 'Brad Pitt is tall' is true or false, it is facts about the meaning of that sentence together with *facts about the height of Brad Pitt*—and it is the latter facts (i.e., the ones I refer to in italics) that are supposed (whose 'job' it is), given the meaning of 'Brad Pitt is tall', to determine whether 'Brad Pitt is tall' is true or false. Thus, in this case, they are the facts of type B, where B is the sentential determination base for 'Brad Pitt is tall'. Now, suppose that one considers all the facts of this type that actually obtain—the B-facts, which in this case are the facts about the height of Brad Pitt—together with the facts about the meaning of 'Brad Pitt is tall' (the M-facts) and one shows that this combination of facts (i.e., the combination of B-facts and M-facts) fails to determine a truth-value for 'Brad Pitt is tall'. Then one has shown that it is indeterminate whether Brad Pitt is tall. Upshot: it is indeterminate whether Brad Pitt is tall iff there is a type of fact B that is a sentential determination base for 'Brad Pitt is tall' but which, in combination with the meaning of 'Brad Pitt is tall', fails to determine a truth-value for the sentence.

#### 4.4.2. *Vagueness-as-absence-of-truth, given the minimal theory*

To return to vagueness-as-absence-of-truth in light of the minimal theory of indeterminacy, here is what supervaluationists are saying when they say that vagueness is a kind of indeterminacy and, accordingly, that there is no fact of the matter about whether (to keep using our example) Brad Pitt is tall. They are saying that, because 'is tall' is vague, the sentence 'Brad Pitt is tall' expresses not a single precise proposition but a set of precise propositions, at least

one of which has a different truth-value than the rest. Supposing, for simplicity's sake, that there are only two (admissible) sharpenings of 'is tall'—say, one on which a male adult is tall iff he is 178 cm tall or above, and another on which a male adult is tall iff he is 182 cm tall or above—the propositions that 'Brad Pitt is tall' expresses are the following:

$p$ : <Brad Pitt is 178 cm tall or above>

$q$ : <Brad Pitt is 182 cm tall or above>

The M-facts are then that 'Brad Pitt is tall' expresses both  $p$  and  $q$ . The B-facts, on the other hand, are that Brad Pitt is 180 cm tall, and so the B-facts determine a truth-value for  $p$  and a truth-value for  $q$ —namely, that  $p$  is true and  $q$  is false. This amounts to the M-facts and the B-facts together failing to determine a truth-value for 'Brad Pitt is tall'. Thus, there is indeterminacy—more specifically, *vagueness-induced* indeterminacy (in that the source of the indeterminacy is the vagueness of 'is tall')—about whether Brad Pitt is tall, and so no fact of the matter.

Another feature of vagueness-as-absence-of-truth that the minimal theory allows us to make clearer (in addition to making clearer the feature that vagueness, on that view, is a kind of indeterminacy) is that vagueness-as-absence-of-truth is consistent with vagueness being a *semantic* phenomenon and with vagueness being a *metaphysical* phenomenon. To begin to see this, consider a rough statement of the distinction—standard in the vagueness literature—between epistemic, semantic and ontic theories of vagueness.

According to epistemic theories of vagueness, vagueness is an epistemic phenomenon, a result of our inability to tell the precise meaning of expressions like 'is tall'. If we were omniscient, there would be no vagueness. By contrast, on semantic theories of vagueness, vagueness is a semantic phenomenon, a result of our linguistic practices not determining a precise meaning for expressions like 'is tall'. If our practices did determine a precise meaning

for such expressions, then, contra epistemic theories, there would be no vagueness. By contrast with both epistemic and semantic theories, some think that even if we were omniscient and our practices did determine a precise meaning for expressions like ‘is tall’, there would still be vagueness. On their ontic theories of vagueness, vagueness is (at least in part) a metaphysical phenomenon—that is, there is vagueness in the world itself, regardless of what we know about the world and of our representations (e.g., our descriptions) of it.<sup>75</sup>

As is obvious, vagueness-as-ignorance is an epistemic theory of vagueness. More important for our purposes here, some versions of vagueness-as-absence-of-truth are *semantic* theories of vagueness while other versions of it are *ontic* theories of vagueness. Why? Because any theory of vagueness that entails truth-value gaps is a version of vagueness-as-absence-of-truth, and there are semantic and ontic theories that entail truth-value gaps.<sup>76</sup> Thus, vagueness-as-absence-of-truth is neutral between (truth-value gap–entailing) semantic and ontic theories of vagueness.<sup>77</sup>

The minimal theory of indeterminacy allows us to make this feature of vagueness-as-absence-of-truth clearer in the following way. Supervaluationism is a version of vagueness-as-absence-of-truth that conceives of vagueness as a semantic phenomenon, and we can make sense of this by appealing to the minimal theory. As we have seen, given supervaluationism, there is (vagueness-induced) indeterminacy about whether Brad Pitt is tall because the combination of the relevant B-facts and M-facts fails to determine a truth-value for ‘Brad Pitt is tall’. However, that combination fails to determine a truth-value because of the M-facts, not because of the B-facts, since the latter facts determine a truth-value for *p* and a truth-value for

---

<sup>75</sup> This way of distinguishing epistemic, semantic and ontic theories is loosely based on Barnes’ [2010] definition of ontic vagueness.

<sup>76</sup> For ontic theories that entail truth-value gaps, see, for instance, Tye [1990], van Inwagen [1990] and Williams [2008].

<sup>77</sup> This is especially significant in light of Schoenfield’s [2016] already mentioned argument that, conditional on moral realism, moral vagueness is a metaphysical phenomenon. For a response to Schoenfield, see Sud [2019], who argues that moral realists can conceive of vagueness as a semantic phenomenon if they appeal to conceptual role semantics.

*q*. In other words, it is the M-facts alone that are to blame for the indeterminacy, since ‘Brad Pitt is tall’ expresses a set of precise propositions rather than a single precise proposition.

Conversely, by appealing to the minimal theory, we can in principle make sense of a version of vagueness-as-absence-of-truth that conceives of vagueness as a (purely) metaphysical phenomenon by showing that, on this view, it is the *B*-facts alone that are to blame for the indeterminacy. To use a different example to illustrate this, suppose that the predicate ‘is bald’ refers to the single property of *having fewer than 846 hairs on one’s scalp*, in which case the sentence ‘Brad Pitt is bald’ expresses a single precise proposition:

*p*<sup>#</sup>: <Brad Pitt has fewer than 846 hairs on his scalp>

Suppose further that Brad Pitt is now ninety-years-old and that it is vague whether he has 846 hairs on his scalp—perhaps he has 845 hairs firmly attached to it but one hair that is on the brink of falling.<sup>78</sup> In such a scenario, it is *the facts about how many hairs Brad Pitt has on his scalp* (the *B*-facts), and not the facts about the meaning of ‘Brad Pitt is bald’ (the *M*-facts), that are to blame for the (vagueness-induced) indeterminacy about whether (ninety-year-old) Brad Pitt is bald.

#### 4.5. *The Up-to-You account of decision-making under symmetric vague betterness*

Assume vagueness-as-absence-of-truth. Then, since Francesca (to start again with her case) knows that it is vague whether *Aborting would be a moral mistake*, she knows that it is neither true nor false that *Aborting would be a moral mistake*. Similarly, since she knows that it is vague whether *Carrying-the-pregnancy-to-term would be a prudential mistake*, she knows that it is neither true nor false that *Carrying-the-pregnancy-to-term would be a prudential mistake*. Which doxastic attitude ought she to have towards these propositions?

---

<sup>78</sup> This is based on an example of ontic vagueness suggested by Barnes [2010: 605].

In Chapter 1, I argued that, together with the Truth Norm of Belief (TNB) (according to which: one may believe  $p$  iff  $p$  is true), supervenience entails Must Not Believe It:

*Must Not Believe It:* If one knows that it is vague whether  $p$ , then one ought not to believe  $p$  and one ought not to believe not- $p$ .

In fact, though, the argument generalises to *any* version of vagueness-as-absence-of-truth: if one knows that it is vague whether  $p$ , then, by vagueness-as-absence-of-truth, one knows that neither  $p$  nor not- $p$  is true, and so, by TNB, one ought not to believe  $p$  and one ought not to believe not- $p$ . Now, Must Not Believe It tells us which doxastic attitudes Francesca ought *not* to have towards *Aborting would be a moral mistake* and *Carrying-the-pregnancy-to-term would be a prudential mistake*: she ought not to *believe* them and she ought not to *disbelieve* them (i.e., she ought not to believe their negations). Thus, what Must Not Believe It does is to narrow down the scope of doxastic attitudes it is (rationally) permissible for her to have. There is a standard doxastic attitude she might have, however, that is consistent with Must Not Believe It, namely, the one that completes the triad of doxastic attitudes (together with belief and disbelief): suspension of judgment (agnosticism).<sup>79</sup>

What is to suspend judgment about whether  $p$ ? It is not simply to fail to believe  $p$  and to fail to disbelieve  $p$ , since we fail to believe and disbelieve many propositions about which we are not suspending judgment (e.g., those we have never considered). As Jane Friedman [2013a, 2017] has argued, to suspend judgment about whether  $p$  is to have an opinion on the relevant matter, not to lack one. More specifically, it is to have a commitment to being neutral or undecided both about whether  $p$  is true and about whether  $p$  is false. Thus, if Francesca

---

<sup>79</sup> Notice that here, as in Chapter 1, I restrict my discussion of doxastic attitudes to those in ‘traditional’ epistemology—hence why I focus on the triad of belief, disbelief and suspension of judgment. In ‘formal’ epistemology, there are many doxastic attitudes one can have towards  $p$ , where these correspond to the many degrees of confidence or belief (credences) one can have in the truth of  $p$ /not- $p$ . For an argument that traditional epistemology’s triad of doxastic attitudes cannot be reduced to credences, see Friedman [2013b].

suspended judgment about whether *Aborting would be a moral mistake* and about whether *Carrying-the-pregnancy-to-term would be a prudential mistake*, she would be in a state of committed neutrality or indecision towards both the truth and the falsity of those propositions.

Is it then permissible for Francesca to suspend judgment (be agnostic) about those propositions? There is strong reason to think not, given vagueness-as-absence-of-truth. To see this, compare her choice situation given vagueness-as-ignorance to her choice situation given vagueness-as-absence-of-truth. In Francesca's situation given vagueness-as-ignorance, it is permissible for her to suspend judgment about those propositions because she knows that there are relevant facts about which she is ignorant—for instance, she knows that there is a fact of the matter about whether Aborting would be a moral mistake but that she is ignorant of it. In her situation given vagueness-as-absence-of-truth, however, she knows that there is *no* relevant fact about which she is ignorant. After all, she knows that both *Aborting would be a moral mistake* and *Carrying-the-pregnancy-to-term would be a prudential mistake* are neither true nor false. Thus, it would be impermissible for her, given what she knows, to inquire further into whether, for instance, Aborting would be a moral mistake—for she knows that there is *no* fact of the matter about whether it would be a moral mistake, and so no fact to be known. Accordingly, there is nothing for her to suspend judgment about.<sup>80</sup> Francesca ought *not* to suspend judgment about either of those propositions.

What follows? That Francesca ought not to have *any* doxastic attitude—belief, disbelief or suspension of judgment—towards those propositions. Thus, no matter what she does, whether she chooses Aborting or Carrying-the-pregnancy-to-term, she ought not to believe that / disbelieve that / suspend judgment about whether she is making a mistake in choice. More specifically, she ought to believe that *Aborting would not be a prudential mistake*, but she ought not to believe that / disbelieve that / suspend judgment about whether *Aborting would*

---

<sup>80</sup> Friedman makes a similar point to mine here in her work on suspension of judgment. See Friedman [2013a: 173, n. 13; 2017: 324, n. 29].

*be a moral mistake*, and she ought to believe that *Carrying-the-pregnancy-to-term would not be a moral mistake*, but she ought not to believe that / disbelieve that / suspend judgment about whether *Carrying-the-pregnancy-to-term would be a prudential mistake*. In other words, she ought to be in a state of mind such that, whatever she does, she is unable to pursue a course of action that ensures that she is (determinately) not making a mistake in choice.

Given that epistemic rationality requires Francesca, due to vagueness, to have no doxastic attitude about whether she would be making a mistake by Aborting and no doxastic attitude about whether she would be making a mistake by Carrying-the-pregnancy-to-term, *epistemic rationality cannot prevent the vagueness in 'better than' from percolating up to the practical realm* (i.e., from 'infecting' what she rationally may and ought to do). Contrast this case—a case of symmetric vague betterness—with cases of *asymmetric* vague betterness, for which an account of decision-making was developed in Chapter 1. In cases of asymmetric vague betterness in the context of choice, option A is (determinately) not worse than option B, and it is vague whether B is not worse than A (equivalently: it is vague whether A is better than B, and B is determinately not better than A). If one knows that one is in that choice situation, epistemic rationality requires us (given vagueness-as-absence-of-truth) to *believe* that one would not be making a mistake by choosing A and to *not* believe that—and indeed to also not disbelieve that or suspend judgment about whether—one would not be making a mistake by choosing B. In that way, in such cases, epistemic rationality *prevents* the vagueness in 'better than' from percolating up to the practical realm, namely, by requiring us to believe that there is exactly one option that, if chosen, would allow us to not make a mistake in choice—hence my account of decision-making under such cases (according to which one rationally ought to choose A).

By contrast, in Francesca's case, there is no way in which epistemic rationality prevents the vagueness in 'better than' from percolating up to the practical realm. Thus, I take it that it does percolate up, generating (vagueness-induced) indeterminacy—and so *no* fact of the matter—about whether practical rationality requires her to choose Aborting and about

whether it requires her not to choose it (i.e., whether it requires her to choose Carrying-the-pregnancy-to-term).

Though I shall not spell this out here, what I have just said about Francesca's case applies, *mutatis mutandis*, to Tomaso's: given what epistemic rationality requires, it is indeterminate whether practical rationality requires him to choose Eating-the-meal and indeterminate whether it requires him not to choose it (i.e., whether it requires him to choose Refraining-from-eating-it).

#### 4.5.1. *The Action Under Normative Indeterminacy Question*

Conditional on vagueness-as-absence-of-truth, then, the right account of decision-making under these cases of symmetric vague betterness is that there is (vagueness-induced) indeterminacy about what one rationally may and ought to do. Because this is (conditional on vagueness-as-absence-of-truth) the right account of decision-making under these cases of symmetric vague betterness, however, there is a further question one might ask:

##### *Action Under Normative Indeterminacy Question:*

What are the features of one's action when one is acting in a choice situation such that it is indeterminate whether one ought to  $\varphi$  and indeterminate whether one ought not to  $\varphi$ ? In other words, what are the features of action under normative indeterminacy?

There are two versions of this question one might consider. On one version, the question is: what are the features of one's action when one is acting—and *knows or at least believes one is acting*—in a choice situation such that it is indeterminate whether one ought to  $\varphi$  and indeterminate whether one ought not to  $\varphi$ ? Let us say that this is action under *suspected* normative indeterminacy (because, at the very least, one believes one is acting under normative indeterminacy). On another version, the question is: what are the features of one's action when one is acting—and *does not know or even believe one is acting*—in a choice situation such that it is



indeterminate whether one ought to  $\varphi$  and indeterminate whether one ought not to  $\varphi$ ? Let us say that this is action under *unsuspected* normative indeterminacy (because one does not even believe one is acting under normative indeterminacy).

In Chapter 3, I will identify what I think are the key features of action under unsuspected normative indeterminacy. For my purposes here, however, it is enough to identify one feature of action under normative indeterminacy that holds regardless of whether the normative indeterminacy is suspected or unsuspected: if one's choice situation is such that it is indeterminate whether one ought to  $\varphi$  and indeterminate whether one ought not to  $\varphi$ , then—to put it in supervenient terms—one is *forced to act in accordance with some*, but in *conflict* with other, (admissible) *sharpenings of the normative indeterminacy*. In other words, one cannot escape acting under some, but against other, sharpenings of the normative indeterminacy. I call this the *inescapability feature* of action under (suspected and unsuspected) normative indeterminacy.

To make this clearer, recall first that, according to supervenientism, truth is truth under all equally acceptable ways of drawing the boundary between the things to which (to focus only on predicates) 'is  $F$ ' applies and those to which it does not—that is, truth is truth under all (admissible) sharpenings. To illustrate this with Francesca's case: if Francesca were at the beginning of the series of Abortions, all the sharpenings of 'is morally permissible' would *agree* that Aborting is morally permissible for her, in which case the proposition *It is morally permissible for Francesca to choose Aborting* would be true under all sharpenings, and so true. Conversely, if she were at the end of the series of Abortions, all the sharpenings of 'is morally permissible' would *disagree* that Aborting is morally permissible for her, in which case *It is morally permissible for Francesca to choose Aborting* would be false under all sharpenings, and so false. But Francesca is neither (determinately) at the beginning of the series of Abortions nor (determinately) at the end of it—she is in the zone of vagueness. In the zone of vagueness, some sharpenings of 'is morally permissible' agree that Aborting is morally permissible for her

while others disagree, which is why *It is morally permissible for Francesca to choose Aborting* is neither true nor false (i.e., it is true on some sharpenings, but false on others).

Now suppose that Francesca (who is in the zone of vagueness) chooses Aborting. Then she is acting in accordance with the sharpenings of ‘is morally permissible’ on which Aborting is morally permissible for her, and so in accordance with the sharpenings on which Aborting is all-things-considered better than Carrying-the-pregnancy-to-term—and in turn in accordance with the sharpenings on which *she rationally ought to choose Aborting*. But then she is not acting in accordance with the sharpenings on which Aborting is *not* morally permissible for her, nor therefore in accordance with the sharpenings on which she rationally ought to choose Carrying-the-pregnancy-to-term. Instead, she is acting in *conflict* with (or *against*) these sharpenings. (*Mutatis mutandis* if she chooses Carrying-the-pregnancy-to-term.)

The moral that I draw from this feature of action under normative indeterminacy is that, since Francesca can only act in accordance with some but not all sharpenings (and no sharpening draws the boundary of ‘is morally permissible’—the source of the indeterminacy—better than any other), it is up to her which sharpenings to act under. In this sense, it is *up to her what to do*. This is of course not to say that it is rationally permissible for her to choose Aborting and rationally permissible for her to choose Carrying-the-pregnancy-to-term, for, as I have just argued, it is *indeterminate* whether it is rationally permissible for her to choose either of those options. Rather, it is to say that there is a sense in which she is *free* to choose Aborting and *free* to choose Carrying-the-pregnancy-to-term, which is distinct from the sense of ‘free’ in which she would be free to choose either were both options rationally permissible. In what sense of ‘free’ then is she free to choose Aborting and free to choose Carrying-the-pregnancy-to-term? In the (special) sense that it is both not *determinately* irrational for her to choose Aborting and not *determinately* irrational for her to choose Carrying-the-pregnancy-to-term. For this reason, I call the account of decision-making developed in this section the *Up-to-You account* of decision-making under symmetric vague betterness.

## 5. IMPLICATIONS FOR DEBATES IN PRACTICAL ETHICS

What are the implications of the two accounts of decision-making developed in Section 4 for the abortion and animal welfare debates? Consider first the implications of the account of decision-making that I argued to be the right account under cases like Francesca's and Tomaso's *if* vagueness-as-ignorance is true—the Play-it-Safe account. The upshot of the Play-it-Safe account for the abortion debate is significant. To see this, recall that, as discussed in Section 2, a view like the simple liberal view holds that, in cases like Francesca's, aborting is morally permissible at the early stage of the pregnancy but not at the late stage of it, and that between these stages there is a zone of vagueness. Given the Play-it-Safe account, a view like the simple liberal view is in practice more *conservative* than has been acknowledged. After all, the view is committed to holding that if a woman finds herself in a choice situation like Francesca's, then it is (determinately) rationally *impermissible* for her to abort both at the late stage of the pregnancy *and* in the zone of vagueness. This implies that only at the early stage of the pregnancy is it (determinately) rationally permissible for her to abort.

What about the implications of the Play-it-Safe account for the animal welfare debate? They are unsurprising. Suppose that Tomaso wants to eat, for example, an oyster-based meal, but that this meal is in the zone of vagueness of the series of Meals (because, let us assume, it is vague whether oysters are sentient—that is, vague whether there is something it is like to be an oyster). Given the Play-it-Safe account, a view like simple moral vegetarianism is committed to holding that if one finds oneself in a choice situation like Tomaso's, then (determinately) one rationally ought *not* to eat that meal. I take this to be unsurprising implication because it is in line with the view of several moral vegetarians, including Singer, who said the following about oysters:

With creatures like oysters, doubts about a capacity for pain are considerable; and in the first edition of [*Animal Liberation*] I suggested that somewhere between a shrimp and an oyster seems as good a place to draw the line as any. Accordingly, I continued occasionally to eat oysters, scallops, and mussels for some time after I became, in every other respect, a vegetarian. But while one cannot with any confidence say that these creatures do feel pain, so one can equally have little confidence in saying that they do not feel pain. Moreover, if they do feel pain, a meal of oysters or mussels would inflict pain on a considerable number of creatures. Since it is so easy to avoid eating them, I now think it better to do so. [Singer 2002: 173-4.]

Consider now the implications of the account of decision-making that I argued to be the right account under cases like Francesca's and Tomaso's *if* vagueness-as-absence-of-truth is true—the Up-to-You account. The upshot of the Up-to-You account for the animal welfare debate is significant. Why? Because it implies that, in cases like Tomaso's, a view like simple moral vegetarianism is committed to holding that there is *no* fact of the matter about whether practical rationality permits eating (to use the same example) an oyster-based meal. Thus, if we find ourselves in a choice situation like Tomaso's—in which we like the taste of oysters and find them tastier than any of the other options available—it is *up to us what to do*.

What about the implications of the Up-to-You account for the abortion debate? They are unsurprising. Given the Up-to-You account, a view like the simple liberal view is committed to holding that, in cases like Francesca's, it is *up to the woman* what to do both at the early stage of the pregnancy and in the zone of vagueness. I take this to be an unsurprising implication because it is in line with the view of many liberals about abortion.

The overall upshot is then that whichever of those two theories of vagueness is true (if any is), there is a significant implication for a debate in practical ethics. If vagueness-as-ignorance is true, there is a significant implication for the abortion debate. If vagueness-as-absence-of-truth is true, there is a significant implication for the animal welfare debate. Notice, however, that these are merely two examples of debates in ethics for which vagueness-as-ignorance and vagueness-as-absence-of-truth have implications, since my argument generalises to any debate in which there are cases structurally identical to Francesca's and Tomaso's—that

is, any debate in which there are cases with, roughly, the following structure:  $\phi$ ing would (determinately) not be a prudential mistake but it is vague whether it would be a moral mistake, and not  $\phi$ ing would (determinately) not be a moral mistake but it is vague whether it would be a prudential mistake.

## 6. THE NEW VAGUENESS VIEW, REVISITED

What is it to defend the new vagueness view of superhard comparisons and choices? On the one hand, it is to defend an account of decision-making under asymmetric vague betterness according to which in any case of this type—that is, in any case in which  $x$  is (determinately) not worse than  $y$ , and it is vague whether  $y$  is not worse than  $x$ <sup>81</sup>—one rationally ought to choose  $x$  over  $y$ . On the other hand, it is to defend an account of decision-making under symmetric vague betterness according to which in at least some cases of this type—that is, in at least some cases in which it is vague whether  $x$  is not worse than  $y$ , and vague whether  $y$  is not worse than  $x$ <sup>82</sup>—it is vague whether one rationally may choose  $x$  over  $y$  and vague whether one rationally may choose  $y$  over  $x$ .

In Chapter 1, I defended the relevant account of decision-making under asymmetric vague betterness, conditional on any Must Not Believe It–entailing theory of vagueness. In Section 4 of this chapter, I defended the relevant account of decision-making under symmetric vague betterness, conditional on any version of vagueness-as-absence-of-truth (which, as we saw, entails Must Not Believe It). This completes my defence of the new vagueness view, conditional on vagueness-as-absence-of-truth. Let us call this the *TVG vagueness view* of superhard comparisons and choices (where ‘TVG’ stands for ‘truth-value gap’).

---

<sup>81</sup> Equivalently: in any case in which it is vague whether  $x$  is better than  $y$ , and that  $y$  is (determinately) not better than  $x$ .

<sup>82</sup> Equivalently: in at least some cases in which it is vague whether  $x$  is better than  $y$ , and vague whether  $y$  is better than  $x$ .

The TVG vagueness view is but one version of the new vagueness view, however. In what follows, I distinguish the TVG vagueness view from other versions of the new vagueness view. In Section 6.1, I distinguish it from an obvious alternative to it, the *IG vagueness view* (where ‘IG’ stands for ‘Ignorance’), which is conditional on vagueness-as-ignorance. In Section 6.2, I distinguish both the TVG and IG vagueness views from another version of the new vagueness view: what I call the *PL vagueness view* (where ‘PL’ stands for ‘pluralist’). In distinguishing the TVG vagueness view from the IG and PL vagueness views, my aim is not to defend any of these two other versions of the new vagueness view, but simply to provide the basic framework for defending them.

### 6.1. *The TVG and IG vagueness views*

Given the discussion in Chapter 1 and in previous sections of this chapter, there is a straightforward way whereby to distinguish different versions of the new vagueness view, namely, by considering their answers to the following question:

*Doxastic Attitude Question for Vague Betterness:*

Which doxastic attitude—belief, disbelief or suspension of judgment—ought one to have towards  $x$  *is not worse than*  $y$  if one knows that it is vague whether  $x$  *is not worse than*  $y$ ?

Before doing so, recall first that in Chapter 1 I argued that, together with the Knowledge Norm of Believe (KNB) (according to which: one may believe  $p$  iff one knows that  $p$ ), Williamson’s version of vagueness-as-ignorance entails Must Not Believe It. In fact, however, the argument generalises to any version of vagueness-as-ignorance (not just Williamson’s): if one knows that it is vague whether  $p$ , then, by vagueness-as-ignorance, one knows that one does not know that  $p$ , and so, by KNB, one ought not to believe  $p$  and one ought not to believe not- $p$ . Let us then assume that vagueness-as-ignorance does entail Must Not Believe It. Given this assumption, both the TVG vagueness view and the IG vagueness view agree on part of the answer to the

Doxastic Attitude Question for Vague Betterness: one ought not to believe *x is not worse than y* and one ought not to disbelieve *x is not worse than y*—an instantiation of Must Not Believe It. However, as the discussion in Section 4 shows, they provide *different* answers to the Doxastic Attitude Question for Vague Betterness, for consider this question: is it permissible to suspend judgment about whether *x is not worse than y* if one knows that it is vague whether *x is not worse than y*? On the TVG vagueness view, the answer is ‘No’. On the IG vagueness view, the answer is ‘Yes’. The upshot is that the TVG vagueness view and the IG vagueness view answer the Doxastic Attitude Question for Vague Betterness, respectively, as follows:

*No Doxastic Attitude:* If one knows that it is vague whether *x is not worse than y*, then one ought not to believe that / disbelieve that / suspend judgment about whether *x is not worse than y*.

*Suspend Judgment:* If one knows that it is vague whether *x is not worse than y*, then one ought to suspend judgment about whether *x is not worse than y*.

In Section 4, I argued for the TVG vagueness view, since I argued that, conditional on vagueness-as-absence-of-truth, there are cases of symmetric vague betterness whose vagueness percolates up, generating vagueness-induced indeterminacy in rational choice. Whether the IG vagueness view can be defended on a similar basis is an open question. In other words, whether there are cases of symmetric vague betterness that, conditional on vagueness-as-ignorance, generate vagueness in rational choice is an open question. Why? Because what I argued in Section 4 is that there are cases of symmetric vague betterness that, conditional on vagueness-as-ignorance, do *not* generate vagueness in rational choice. Thus, anyone who wishes to argue for the IG vagueness view would have to show that only some, not all, cases of symmetric vague betterness generate vagueness in rational choice.

## 6.2. The PL vagueness view

According to the versions of the new vagueness view that I have introduced so far, there is a straightforward answer to the Doxastic Attitude Question for Vague Betterness. On the TVG vagueness view, that answer is No Doxastic Attitude: one ought not to believe that / disbelieve that / suspend judgment about whether *x is not worse than y*. On the IG vagueness view, that answer is Suspend Judgment: one ought to suspend judgment about whether *x is not worse than y*.

According to the alternative version of the new vagueness view that I introduce here, the answer is: it depends. To spell it out, notice that knowing that it is vague whether *x is not worse than y* does not by itself tell us whether one is before a case of *symmetric* vague betterness or before a case of *asymmetric* vague betterness. To know whether one is before a symmetric vague betterness case or before an asymmetric vague betterness case, one would also have to know the status of *y is not worse than x*. The thought is then this. The answer to the Doxastic Attitude Question for Vague Betterness depends on the known status of *y is not worse than x*—whether it is known to be true or vague. If one knows that *y is not worse than x*, then the answer to it is, let us say, A1. If one knows that it is vague whether *y is not worse than x*, then the answer to it is, let us say, A2. To put it another way, different *types* of vague betterness cases are governed by *different doxastic norms*, with any case of the asymmetric vague betterness-type being governed by the norm underlying A1 and any case of the symmetric vague betterness-type being governed by the norm underlying A2. This is what I call the *PL vagueness view* of superhard comparisons and choices.

Can we spell out the PL vagueness view further, specifying the doxastic norm governing any case of asymmetric vague betterness and the doxastic norm governing any case of symmetric vague betterness? To see that we can, note, to begin with, that one way of understanding my main argument in Chapter 1—the Make No Mistake argument—is that the doxastic norm governing any case of asymmetric vague betterness is Must Not Believe It, since



it is the doxastic norm that allows us to develop the right account of decision-making under asymmetric vague betterness. On the PL vagueness view, however, Must Not Believe It should be understood in a restricted way, as applying only to some propositions:

*Must Not Believe It, restricted version:* For some proposition  $p$ , if one knows that it is vague whether  $p$ , then one ought not to believe  $p$  and one ought not to believe not- $p$ .

According to the PL vagueness view, in any case in which one knows that it is vague whether  $x$  is not worse than  $y$ , and in which one knows that  $y$  is not worse than  $x$  (i.e., whenever one knows that one is before a case of asymmetric vague betterness), the restricted version of Must Not Believe It applies to  $x$  is not worse than  $y$ , entailing that one ought not to believe or disbelieve  $x$  is not worse than  $y$ . For cases of asymmetric vague betterness in the context of choice, this entails—together with Make No Mistake—my account of decision-making under asymmetric vague betterness.

What is the doxastic norm governing any case of symmetric vague betterness, according to the PL vagueness view? In Chapter 1, I considered a theory of vagueness—non-standard supervaluationism<sup>83</sup>—which (together with TNB) entails the following doxastic norm:

*Vague Whether One Must Not Believe It:* If one knows that it is vague whether  $p$ , then it is vague whether one ought not to believe  $p$  and vague whether one ought not to believe not- $p$ .

Recall why: if one knows that it is vague whether  $p$ , then, by non-standard supervaluationism, one knows that it is vague whether  $p$  is true or false (i.e., one knows that it is vague which

---

<sup>83</sup> Just like supervaluationism is a version of what we have been calling vagueness-as-absence-of-truth, non-standard supervaluationism is a version of what one might call vagueness-as-indeterminate-truth—a type of theory of vagueness that is neutral between (indeterminate truth-entailing) semantic and ontic theories of vagueness (e.g., Barnes' [2010] ontic theory is a version of it).

truth-value  $p$  has, true or false). If one knows that it is vague whether  $p$  is true or false, then, by TNB, Vague Whether One Must Not Believe It follows.<sup>84</sup> As I argued in Chapter 1, Vague Whether One Must Not Believe It yields the *wrong* account of decision-making under asymmetric vague betterness and should for that reason be rejected. However, that result is consistent with a restricted version of it being true:

*Vague Whether One Must Not Believe It, restricted version:* For some proposition  $p$ , if one knows that it is vague whether  $p$ , then it is vague whether one ought not to believe  $p$  and vague whether one ought not to believe not- $p$ .

According to the PL vagueness view, in any case in which one knows that it is vague whether  $x$  is not worse than  $y$ , and in which one knows that it is vague whether  $y$  is not worse than  $x$  (i.e., whenever one knows that one is before a case of symmetric vague betterness), the restricted version of Vague Whether One Must Not Believe It applies to both  $x$  is not worse than  $y$  and  $y$  is not worse than  $x$ , entailing that it is vague whether one ought not to believe or disbelieve  $x$  is not worse than  $y$  and vague whether one ought not to believe or disbelieve  $y$  is not worse than  $x$ . The upshot is that there is vagueness about which option,  $x$  or  $y$ , to choose. This yields an account of decision-making under symmetric vague betterness according to which there is vagueness in rational choice in any vague betterness case of this type.

Thus, on the PL vagueness view, any case of asymmetric vague betterness is governed by (the restricted version of) Must Not Believe It, and any case of symmetric vague betterness is governed by (the restricted version of) Vague Whether One Must Not Believe It. Accordingly, the PL vagueness view provides two answers to the Doxastic Attitude Question for Vague Betterness, depending on the known status of  $y$  is not worse than  $x$ :

---

<sup>84</sup> As should be clear, this argument generalises to *any* version of vagueness-as-indeterminate-truth.

*A1: If one knows that it is vague whether  $x$  is not worse than  $y$ , and one knows that  $y$  is not worse than  $x$ , then one ought not to believe or disbelieve  $x$  is not worse than  $y$ , and*

*A2: If one knows that it is vague whether  $x$  is not worse than  $y$ , and one knows that it is vague whether  $y$  is not worse than  $x$ , then it is vague whether one ought not to believe or disbelieve  $x$  is not worse than  $y$ , and vague whether one ought not to believe or disbelieve  $y$  is not worse than  $x$ ,*

where, to repeat, the doxastic norm underlying A1 is (the restricted version of) Must Not Believe It while the doxastic norm underlying A2 is (the restricted version of) Vague Whether One Must Not Believe It.

Whether the PL vagueness view is defensible is an open question. But one way to see the ‘big-picture’ motivation for it would be by considering an analogy with contingency. Consider the following question:

*Doxastic Attitude Question for Contingency:*

Which doxastic attitude—belief, disbelief or suspension of judgment—ought one to have towards  $p$  if one knows that it is contingent whether  $p$ ?

The answer to this question is: it depends. In some cases, one ought to believe  $p$ . In others, one ought to disbelieve  $p$ . In still others, one ought to suspend judgment about whether  $p$ . It all depends on one’s evidence. In other words, there is no single doxastic attitude one ought to have towards propositions that one knows to be contingent. Perhaps vagueness (and indeterminacy, more generally) is like contingency in this way. If so, then the answer to the analogous question for vagueness,

*Doxastic Attitude Question for Vagueness:*

Which doxastic attitude—belief, disbelief or suspension of judgment—ought one to have towards  $p$  if one knows that it is vague whether  $p$ ?

would be that there is no single doxastic attitude one ought to have towards propositions that one knows to be vague. Different cases of vagueness require different doxastic attitudes.<sup>85</sup> As should be clear by now, the PL vagueness view relies on a view along these lines and applies it to vague betterness cases, holding that the two types of vague betterness are each associated with a different doxastic norm.

## 7. CONCLUSION

In this chapter, I linked the literature on superhard comparisons and choices with the literature on moral vagueness, with the aim of advancing both. I began by showing that there are cases of moral vagueness that, properly understood, turn out to also be cases of symmetric vague betterness. Then I developed two distinct accounts of decision-making under these cases of symmetric vague betterness, one conditional on vagueness-as-ignorance and another on vagueness-as-absence-of-truth. The aim was twofold. On the one hand, I aimed to show that these accounts of decision-making have distinct implications for debates in ethics whose cases are structurally identical to Francesca's and Tomaso's. On the other hand, I aimed to complete my defence of the new vagueness view of superhard comparisons and choices, conditional on vagueness-as-absence-of-truth.

---

<sup>85</sup> A view along these lines is Williams' normative silence view of indeterminacy. See Williams [2012] for a defence, and Williams [2014b: esp. 420-1] for an application of it to cases of indeterminacy in personal identity. For a critique of Williams' view, see Eklund [2013].

## *Chapter 3*

### Superhard choices and indeterminate art

Can one make works which are not works of ‘art’?

—Marcel Duchamp<sup>86</sup>

#### 1. INTRODUCTION

In January 1917, the first annual exhibition of the American Society of Independent Artists was announced. The Society had been founded recently, in December 1916, by the art collector Walter Arensberg and several artists, among whom was the French artist Marcel Duchamp. It had a clear aim, as per the announcement: ‘The Society of Independent Artists has been incorporated under the laws of New York for the purpose of holding exhibitions in which all artists may participate independently of the decisions of juries.’<sup>87</sup> In response to the announcement, Duchamp—who was a member of the board of directors—submitted an entry, together with a payment of six dollars, under the pseudonym ‘R. Mutt’. The six-dollar

---

<sup>86</sup> Handwritten note from 1913 in ‘A L’Infinifit’ (also known as ‘The White Box’). See Duchamp [1912-20: 74].

<sup>87</sup> See Glackens et al. [1917: fol. 3].

payment was all that was needed to join the Society, membership of which entitled anyone, according to the by-laws of the Society, to show their work at the annual exhibition. Duchamp's entry was an ordinary, store-bought urinal titled *Fountain*, which he signed and dated. By the time *Fountain* arrived, an argument broke out between some members of the board of directors, as was later recounted by Beatrice Wood, an artist and friend of Duchamp with whom he edited the magazine *The Blind Man*. 'This is what the whole exhibit is about; an opportunity to allow an artist to send in anything he chooses, for *the artist* to decide what is art, not someone else', remarked Arensberg (who had been an accomplice of Duchamp in the whole affair), to which the American artist George Bellows retorted: 'You mean to say, if a man sent in horse manure glued to a canvas that we would have to accept it?' [Wood 1985: 29-30]. In the end, a majority vote was taken among the members of the board present on the occasion, the outcome of which is well known: *Fountain* was not exhibited.<sup>88</sup> In protest, Duchamp, never revealing that he himself had sent the urinal, resigned from the board.<sup>89</sup>

Widely regarded by art historians and critics as one of the most influential works—if not the most influential work—of twentieth-century art, *Fountain* raises several philosophical problems. Two of these problems have been the focus of much discussion. One of them is a theoretical problem: is *Fountain* art?<sup>90</sup> The other is a practical problem: how, if at all, should

---

<sup>88</sup> What is less well known is that no scandal occurred following the rejection of *Fountain*. The few newspapers that picked up on the affair did not publish any photograph of *Fountain*, merely describing it as a 'bathroom fixture' and not looking further into the matter (see Camfield [1989: 27-8]). In fact, the only significant press report was published in *The Blind Man*, the obscure, avant-garde magazine co-edited by Duchamp himself in which the now famous Alfred Stieglitz photo of *Fountain* appeared. As the art historian Elena Filipovic [2016: 83] recently put it, 'the rejection of *Fountain* was no more than a footnote to a small, local story'. I return to the set of events that led to the rejection of *Fountain* several times throughout this chapter, starting in Section 2.

<sup>89</sup> The urinal that Duchamp originally submitted to the annual exhibition disappeared soon after the events narrated here and it is not clear what happened to it. The urinals bearing the name of *Fountain* that are now on display in leading art museums, such as Tate Modern in London and the Moderna Museet in Stockholm, are versions of the work that Duchamp authorised in the early 1960s, with the exception of the one on display at the Philadelphia Museum of Art, which dates from the 1950s. None of this matters a great deal for my purposes in what follows, but it is revealing, as we shall see, that only in the 1950s–60s did new versions of *Fountain* start to circulate.

<sup>90</sup> This problem is widely discussed in the context of a more general problem, the problem of what art is. See especially Danto [1964, 1981], Dickie [1969, 1974, 1984], Beardsley [1983], Levinson [1979, 1989, 1993], Stecker [1997], Davies [1991], Gaut [2000], Zangwill [2007], Abell [2012], and Lopes [2008, 2014].

*Fountain* be appreciated as art?<sup>91</sup> Though the theoretical problem and the practical problem (as I shall call them, respectively) tend to be discussed in isolation, there is much that they have in common. One, often neglected, thing that they have in common is that for each of them a corresponding *meta*-problem arises. For the theoretical problem, the theoretical *meta*-problem (as I shall call it) arises: why does *Fountain* pose theoretical puzzlement, namely, puzzlement about whether it is art?<sup>92</sup> For the practical problem, the practical *meta*-problem (as I shall call it) arises: why does *Fountain* pose practical puzzlement, namely, puzzlement about how, if at all, it should be appreciated as art?<sup>93</sup>

In an important sense, art historians need not be concerned about what the correct solution to the theoretical problem turns out to be. Art historians are in the business of studying things that are *treated as art*, and it is an empirical fact that *Fountain* is treated as art. That certain things, but not others, are treated as art is all that art historians need to formulate two problems in which *Fountain* plays a fundamental role. One of these problems arises from the empirical observation that, sometime after Duchamp's submission of *Fountain*, there was a radical change in the *practice of treating things as art* (as I shall refer to what is most commonly referred to as the practice of art). After all, before *Fountain* and Duchamp's other 'readymades'—the name that he himself gave the everyday objects which he signed, dated and titled—ordinary urinals, snow shovels, bottle racks, combs and the like were not treated as art.<sup>94</sup> This raises a question: when and how did this change in the practice *exactly* occur?<sup>95</sup> As I understand the question, the 'when' requires a (fully) historical explanation while the 'how'

---

<sup>91</sup> This problem is widely discussed in the context of a more general problem, the problem of how, if at all, works of conceptualism—the genre which Duchamp pioneered (especially) through *Fountain*—should be appreciated as art. See especially Shelley [2003], Danto [1981], Binkley [1970], Carroll [1986, 2004], and Costello [2013].

<sup>92</sup> See especially Lopes [2008; 2014: e.g., 71-4].

<sup>93</sup> See especially Hopkins [2007], Lopes [2007], and Davies [2003: e.g., 189-91].

<sup>94</sup> Throughout, I use 'readymade' to refer only to what Duchamp called the *unassisted* readymade, of which *Fountain*, *In Advance of the Broken Arm* (a snow shovel), *Bottle Rack* and *Comb* are examples.

<sup>95</sup> See especially Thierry De Duve's [2013a, 2013b, 2014a, 2014b, 2014c, 2014d] recent series of six essays in the magazine *Artforum*, though my formulation of the problem differs considerably from his, since he only focuses on the 'when' side of the question and does not distinguish being art from being treated as art. See also De Duve [1996].

requires a (partially) philosophical explanation, an account of the mechanisms whereby the change occurred. Call this the *problem of the change in the practice*.

The other problem arises from the empirical observation that the avant-garde of the mid-1950s through the mid-1970s—the so-called *neo-avant-garde*—was in a significant way a reaction to the avant-garde of the early 1910s through the mid-1920s—the so-called *historical avant-garde*. This raises a question: what is the relationship between these two avant-gardes? In particular, what did the neo-avant-garde do that had not already been done, to a large extent, by the historical avant-garde? Call this the *problem of the two avant-gardes*.<sup>96</sup> *Fountain* plays a fundamental role in this problem through the most influential of its sub-problems: what did the neo-avant-garde do—through Conceptualism<sup>97,98</sup>, one of the movements of the neo-avant-garde—that had not already been done, to a large extent, by the historical avant-garde *through Dada*<sup>99</sup> (the movement of the historical avant-garde within which *Fountain* emerged)? Call this the problem of the two avant-gardes *via Dada*.

---

<sup>96</sup> See especially Bürger [1984], Buchloh [1984, 1986, 2000] and Foster [1996].

<sup>97</sup> Two related points on terminology. Throughout the chapter, I use ‘Conceptualism’ rather than the more common ‘Conceptual art’ to remain as neutral as possible on some of the issues that will occupy us in this chapter. Moreover, I use ‘Conceptualism’ (capitalised) to refer to the movement of the neo-avant-garde and ‘conceptualism’ (uncapitalised) to refer to the genre that Duchamp pioneered. Given this, it is correct to say that *Fountain* is a work of conceptualism (or a conceptual work) but not to say that it is a work of Conceptualism (or a Conceptual work).

<sup>98</sup> For those unfamiliar with Conceptualism, here are five typical features of conceptualism (the genre which Duchamp pioneered), outlined by Peter Goldie and Elisabeth Schellekens [2007: xvii–xviii], that are also features of Conceptualism (the movement of the neo-avant-garde): ‘1. [Conceptualism] aims to remove the traditional emphasis on sensory pleasure and beauty, replacing it with an emphasis on ideas and the view that the art object is to be ‘dematerialized’. 2. [Conceptualism] sets out to challenge the limits of the identity and definition of artworks and questions the role of agency in art-making. 3. [Conceptualism] seeks, often as a response to modernism, to revise the role of art and its critics so that art-making becomes a kind of art criticism, at times also promoting anti-consumerist and anti-establishment views. 4. [Conceptualism] rejects traditional artistic media, particularly the so-called plastic arts, in favour of new media of production such as photography, film, events, bodies, mixed media, ready-mades, and more. 5. [Conceptualism] replaces illustrative representation by what some call ‘semantic representation’—semantic not only (not necessarily) in the sense of words appearing on or in the work of art itself, but in the sense of depending on meaning being conveyed through a text or supporting discourse.’

<sup>99</sup> For those unfamiliar with Dada, here is a brief explanation of some its basic tenets, and the context in which it emerged, by the art historian and Dada scholar David Hopkins [2004: xiv]: ‘Dada, born in 1916 and over by the early 1920s, was an international artistic phenomenon, which sought to overturn traditional bourgeois notions of art. It was often defiantly anti-art. More than anything, its participants, figures such as Marcel Duchamp, Francis Picabia, Tristan Tzara, Hans Arp, Kurt Schwitters, and Raoul Hausmann, counterposed their love of paradox and effrontery to the insanities of a world-gone-mad, as the First World War raged in Europe.’



We have then six problems, not all of which are philosophical: the theoretical problem and its meta-problem, the practical problem and its meta-problem, the problem of the change in the practice and the problem of the two avant-gardes via Dada. Each of these problems has *Fountain* at its core. Nevertheless, the current solutions to them are solutions that consider them largely in isolation from each other, as problems to be solved more or less, if not entirely, independently of one another. Is there, however, a *unified* solution to these problems, a solution which takes on them together and solves them all at once? Finding a unified solution is no doubt harder than finding independent solutions, but its consequences would be far-reaching, allowing us, for instance, to integrate the philosophical literature on the theoretical and practical problems with the art historical literature on *Fountain* and the rise of the avant-gardes. Call the harder problem of finding a unified solution to the six (later, seven) problems the *Duchamp's box problem*—so called because Duchamp's act of submitting *Fountain*, like Pandora's act of opening the box in classical Greek mythology, is the source of many problems, and because only through *Box-in-a-Valise* (a kind of portable retrospective exhibition of his work in which he included a miniature replica of *Fountain*) did he officially acknowledge to be behind the Mutt affair.

How does finding a solution to the Duchamp's box problem relate to superhard choices? Not in an obvious way. But to anticipate notice that, in light of the connection between choice and indeterminacy explored in Chapter 2, there is a way in which they might be related. One might think that, when the board of directors of the Society of Independent Artists received *Fountain*, they faced a superhard choice, understood here as a choice situation in which it was indeterminate whether they were required to exhibit *Fountain* at the annual show and indeterminate whether they were required not to exhibit it—in a nutshell, a situation in which there was indeterminacy about what to do, or (more simply put) normative indeterminacy. In fact, this is (very roughly) the hypothesis that I explore in this chapter. In doing so, I develop a framework for a solution to the Duchamp's box problem, the *normative*

*indeterminacy hypothesis*—so called because normative indeterminacy and action under it play a key role in the solution.

The structure of the chapter is as follows. In Section 2, I discuss the problems that make up the Duchamp's box problem in greater detail, and I argue that we need a unified solution to those problems because there is a problem, which has gone wholly unnoticed in the literature, that links them all up—a problem which I call *Duchamp's paradox*. (For this reason, the Duchamp's box problem is made up of seven problems, not just six.) In Section 3, I introduce the bare-bones of the normative indeterminacy hypothesis as a non-standard solution to Duchamp's paradox. In Section 4, drawing in part on Ásta's [2018] recent account of social properties, I flesh out the normative indeterminacy hypothesis. In Section 5, I return to a question left hanging in Chapter 2—the Action Under Normative Indeterminacy Question—and spell out five features of action under (unsuspected) normative indeterminacy, partially on the basis of which I argue that the normative indeterminacy hypothesis yields a solution to the Duchamp's box problem. In Section 6, I consider an alternative interpretation of the normative indeterminacy hypothesis.

## 2. THE DUCHAMP'S BOX PROBLEM

In this section, I spell out the Duchamp's box problem. The plan is as follows. In Section 2.1, I introduce Duchamp's paradox. In Section 2.2, I spell out the theoretical problem and its meta-problem. In Section 2.3, I spell out the two problems from the art historical literature, the problem of the change in the practice and the problem of the two avant-gardes via Dada. In Section 2.4, I spell out the practical problem and its meta-problem. In Section 2.5, I explain why we need a unified solution to all these problems—or, to put it another way, a solution to the Duchamp's box problem.

## 2.1. *Duchamp's paradox*

Let us return to where we started, the Mutt affair. Shortly after the opening of the annual exhibition of the Society of Independent Artists—which displayed a staggering 2500 works by 1200 artists<sup>100</sup>—the board of directors issued a statement about the Mutt affair to the press: ‘The *Fountain* may be a very useful object in its place, but its place is not an art exhibition and it is by no definition a work of art.’<sup>101</sup> That the board of directors (as a collective) believed that *Fountain* is not art should be obvious even in the absence of this statement, since their behaviour—in this case, their decision to not exhibit *Fountain*—cannot plausibly be explained without postulating such a belief.

Notice, however, that that belief *alone* cannot explain the behaviour of the board of directors. After all, the belief that *Fountain* is not art is compatible with it being okay to exhibit *Fountain* at the annual show. To explain their behaviour, then, we need to postulate a further belief, namely, that the board believed that that there is a *norm* governing their practice: *Do not treat as art that which is not art*—or ‘Do Not TAA’ for short. Slightly more formally:

*Do Not TAA*: (For any  $x$ ) One must not: treat  $x$  as art if  $x$  is not art.<sup>102</sup>

To exhibit *Fountain* at the annual show would be to treat it as art. So, together with their belief that *Fountain* is not art, the board’s belief that Do Not TAA is in force explains their behaviour.

(What exactly is it to treat something as art? I take it that we have an intuitive grasp on this, given our prior grasp on what it is to treat  $x$  as  $F$ , what works of art are and what behaviours are appropriate in their presence. Nevertheless, I think that we can make the notion precise. Consider this question: what is the point, or at least one of the points, of exhibiting a work at the Society’s annual show, or for that matter at Tate Modern, the Philadelphia Museum

---

<sup>100</sup> See Camfield [1989: 20].

<sup>101</sup> Quoted in Naumann [2012: 72].

<sup>102</sup> In other words: (for any  $x$ ) if  $x$  is not art, then one ought not to treat  $x$  as art.

of Art, and so on? That it be appreciated as art, where to appreciate something as art is to appreciate what one believes to be the artistically-relevant properties of that thing.<sup>103</sup> Based on this simple observation, we can make the notion precise, as follows:  $x$  is treated as art iff  $x$  is treated as something to be appreciated as art. So, if a work was exhibited at the Society's annual show, then that work would be treated as something to be appreciated as art, and so treated as art.)

Now, as we saw in Section 1, the aim of the Society of Independent Artists, according to the announcement of its first exhibition, was to hold annual exhibitions in which artists could take part without their submitted work having to go through a process of selection by a jury. Moreover, we also saw there that, as the announcement states, there were 'no requirements for admission to the Society save the acceptance of its principles and the payment of the initiation fee of one dollar and the annual dues of five dollars.'<sup>104</sup> Because the response to the announcement was overwhelming, the foreword to the catalogue of the exhibition concluded that this 'proved the need of a society which enables all artists, irrespective of school or group, to exhibit their works in America with complete freedom', going on to state the following:

The governing principle of the Society, which *permits a member to exhibit whatever he wishes* on the payment of nominal dues, is applied to the catalogue of the Exhibition, in which each exhibitor has the *right* to buy the space for one illustration. The illustrations are thus decided upon by no process of selection, but by the spontaneous desire of the exhibitors themselves. [Glackens et al. 1917: fol. 3, my italics]

The reason this is important is that—since the 'governing principle' and the criteria to become a member were an integral part of the Society's by-laws—the members of the Society had the *right*, given the by-laws, to exhibit their submitted work at the annual show *without interference*

---

<sup>103</sup> An artistically-relevant property might be understood as a property to which one can appeal in attributions of value to that thing as art.

<sup>104</sup> Glackens et al. [1917: fol. 4].

(e.g., from a jury). Consequently, if the board of directors interfered with a member's right to exhibit their submitted work—namely, by not exhibiting it at the annual show—then that member would be in a position to demand of the board that they exhibited the work. For this reason, a member's right to exhibit their submitted work is a *demand-right*, as Margaret Gilbert [2018: 70] recently called this kind of right:  $A$  has a demand-right against  $B$  to  $B$ 's  $\varphi$ ing iff  $B$  is obligated to  $A$  to  $\varphi$ .<sup>105</sup> In the context of the Society's annual exhibition, this translates as follows: any member of the Society who submitted their work,  $x$ , to the annual show had a demand-right against the board to their exhibiting  $x$  iff the board had an obligation to that member to exhibit  $x$ .

What Mutt did when he submitted *Fountain*, together with the six dollars, was then to exercise the right that any member of the Society had. In doing so, he put the board of directors in a situation in which it was required for them to exhibit *Fountain*, and thereby required—and so *permissible*—to treat *Fountain* as art. However, together with Do Not TAA, to which the board were also committed, this led to a paradox (which I name after Duchamp because I believe it was he who first discovered it):

*Duchamp's paradox*

- (1) (For any  $x$ ) One must not: treat  $x$  as art if  $x$  is not art. (Do Not TAA)
- (2) *Fountain* may be treated as art (i.e., it is permissible to treat it as art).

Therefore,

- (C) *Fountain* is art.

Why would this argument be regarded as a paradox in 1917 (were those participating in the practice of treating things as art aware of it as formulated above)? Because it appeared to be valid, each of its premises appeared to be true when considered on its own (i.e., independently

---

<sup>105</sup> As those familiar with the literature on rights will no doubt notice, a demand-right is similar to a claim-right on the standard, Hohfeldian analysis of rights, but whether (and, if so, how) they differ from each other does not matter for my purposes here.

of the paradox), and its conclusion appeared to be false when considered on its own. The board's (implicit) reaction to the paradox consisted in denying one of its premises—namely, premise (2)—by way of the following argument:

*The board's reaction to Duchamp's paradox*

(For any  $x$ ) One must not: treat  $x$  as art if  $x$  is not art. (Do Not TAA)

*Fountain* is not art.

Therefore,

*Fountain* must not be treated as art (i.e., premise (2) is false).

Today, some would not regard what I am calling Duchamp's paradox as a paradox, because, to them, (C)—the conclusion of the paradox—no longer appears to be false; on the contrary, their intuitions tell them that (C) is *true*. There is a straightforward explanation for this, however. Although Duchamp's paradox has gone unnoticed in the literature, the theoretical problem (to be discussed below, in Section 2.2) has been the subject of an overwhelming number of reactions and the dominant solution to the latter is that *Fountain* is art. As a result, endorsing the dominant solution to the theoretical problem amounts to endorsing the view that (C), despite appearances to the contrary in 1917, is true. It is unsurprising then that, to some, (C) does not appear to be false any more, since our intuitions have been shaped by the many (implicit) conclusion-endorsing reactions to the paradox, both in the practice and in art theory, since 1917.<sup>106</sup>

---

<sup>106</sup> That said, it is sometimes overlooked that some Dadaists (i.e., some members of Dada) did not endorse (C). Here is what the Dadaist Hans Richter [1965: 90], in his first-hand account of the history of Dada, said about *Fountain*: 'Of course, the bottle-rack and the urinal are not art. But the laughter that underlies this shameless exposure of 'all that is holy' goes so deep that a kind of topsy-turvy admiration sets in which applauds at its own funeral (the funeral, that is, of 'all that is holy').'

## 2.2. *The theoretical problem and its meta-problem*

Is *Fountain* art? This is the theoretical problem. The dominant solution to it, as just mentioned, is that *Fountain* is art.<sup>107</sup> The most prominent representative of this solution is George Dickie's. Dickie's solution is derived from his (version of the) institutional theory of art, according to which (for any  $x$ )  $x$  is a work of art iff  $x$  is an artefact onto which at least one person acting on behalf of a social institution—the artworld—conferred the status of candidate for appreciation [Dickie 1971: 101; 1974: 34].<sup>108</sup> Typically, for any given work of art, the person who, acting on behalf of the artworld, conferred the relevant status onto the work is the artist. As Dickie himself wrote:

[A] number of persons are required to make up the social institution of the artworld, but only one person is required to act on behalf of the artworld and to confer the status of candidate for appreciation. In fact, many works of art are only seen by one person—the one who creates them—but they are still art. The status in question may be acquired by a single person's acting on behalf of the artworld and *treating an artifact as a candidate for appreciation*. Of course nothing prevents a group of persons from conferring the status, but it is usually conferred by a single person, the artist who creates the artifact. [Dickie 1974: 37-8]

According to Dickie, the case of *Fountain* is no exception.<sup>109</sup> Thus, on his solution to the theoretical problem, *Fountain* is art because someone acting on behalf of the artworld—in this case, and as is typical, the artist himself (Duchamp)—conferred onto it the status of candidate for appreciation.

---

<sup>107</sup> The dominant solution is explicitly endorsed by Danto [1964, 1981], Dickie [1974, 1984], Levinson [1979, 1989, 1993], Abell [2012], Lopes [2008, 2014], and many others.

<sup>108</sup> This is Dickie's earlier version of the theory, as he himself now calls it (see Dickie 2000: 93). For the later version, see Dickie [1984]. The differences between these versions do not matter for my purposes here.

<sup>109</sup> Dickie [1974: 38] makes this clear (i.e., that, on his view, it was Duchamp who conferred the relevant status onto *Fountain*) when he compares the case of Duchamp's submission of *Fountain* to the annual exhibition of the Society of Independent Artists with the case of a salesperson's placing plumbing supplies before us: 'It may be helpful to compare and contrast the notion of conferring the status of candidate for appreciation with a case in which something is simply presented for appreciation [...]. 'Placing before' and 'conferring the status of candidate for appreciation' are very different notions, and this difference can be brought out by comparing the salesman's action with the superficially similar act of Duchamp in entering a urinal which he christened *Fountain* in that now-famous art show. The difference is that Duchamp's act took place within the institutional setting of the artworld and the plumbing salesman's action took place outside it. The salesman could do what Duchamp did, that is, convert the urinal into a work of art, but such a thing probably would not occur to him.'

The obvious alternative solution to the theoretical problem is that *Fountain* is not art. The most prominent representative of this solution is Monroe Beardsley's. Beardsley's solution is derived from his (version of the) aesthetic theory of art, according to which (for any  $x$ )  $x$  is a work of art iff  $x$  is created with the intention that it be capable of satisfying the aesthetic interest [Beardsley 1983: 58]. On this solution, *Fountain* is not art because it was not created with the relevant intention. Indeed, as Duchamp made clear, it was created precisely with the intention that it *not* be capable of satisfying the aesthetic interest.<sup>110</sup> At this point, it is worth quoting at length what Beardsley himself wrote about *Fountain*:

[C]ommon sense should not be abandoned along with philosophical acumen in these matters. The fuss about Duchamp's *Fountain* has long amazed me. It does not seem that in submitting that object to the art show [...], Duchamp or anyone else thought of it either as art or as having an aesthetic capacity. He did not establish a new meaning of 'artwork', nor did he really inaugurate a tradition that led to the acceptance of plumbing figures (or other 'readymades') as artworks today. If there was a point, it was surely to prove the jury that even their tolerance had limits, and that they would *not* accept anything—at least gracefully. This small point was made effectively, but the episode doesn't seem to me to provide the slightest reason to regard the aesthetic definition as inadequate. Many objects exhibited today by the avant-garde evidently do make comments of some kind on art itself, but these objects may or may not be artworks. [...] To classify them as artworks because they are called art by those who are called artists because they make things they call art is not to classify at all, but to think in circles. Perhaps these objects deserve a special name, but not the name of art. The distinction between objects that do and those that do not enter into artistic activities by reason of their connection to the aesthetic interest is still vital to preserve, and no other word than 'art' is suitable to make it. [Beardsley 1983: 60]

One challenge that any solution to the theoretical problem faces is that of explaining why *Fountain* poses the theoretical problem in the first place—what I have called the theoretical *meta*-problem. In fact, given how closely linked the theoretical problem and its meta-problem are, we should expect a solution to either of them to yield, or at least provide the basis for, a solution to the other: if we had a solution to the theoretical problem, we would expect it to

---

<sup>110</sup> See Duchamp [1961: 141]: 'A point which I want very much to establish is that the choice of these 'readymades' was never dictated by an esthetic delectation. This choice was based on a reaction of visual indifference with at the same time a total absence of good or bad taste... in fact a complete anesthesia.'



yield, or at least provide the basis for, a solution to the meta-problem; if we had a solution to the meta-problem, we would expect it to yield, or at least provide the basis for, a solution to the theoretical problem. In other words, we should expect for there to be a *unified* solution to the theoretical problem and its meta-problem.

One philosopher who is aware of the need to provide a unified solution to the theoretical problem and its meta-problem is Dominic Lopes [2008, 2014], who has proposed a unified solution to them based on the buck-passing theory of art, according to which (for any  $x$ )  $x$  is a work of art iff, for some kind  $K$ ,  $K$  is an art and  $x$  is a work of  $K$ . On the buck-passing theory, Jacques-Louis David's *The Death of Socrates*, for example, is art because painting is one of the arts—along with sculpture, music, literature, architecture, and so on—and *The Death of Socrates* is a work of painting.

Lopes' solution to the theoretical problem is derived from his particular version of the buck-passing theory, according to which *Fountain* is art because, and despite appearances to the contrary, it belongs to one of the arts. Which one? The *new* art which *Fountain* pioneered: conceptualism (or, as Lopes calls it, conceptual art)—a surprising and controversial claim, since conceptualism is standardly understood not as an art but as a genre, along with abstract art, tragedy, comedy, science fiction, and so on.<sup>111</sup> However, on Lopes' view, it is crucial that *Fountain* belongs to one of the arts *despite appearances to the contrary* and that conceptualism is one of the arts *despite appearances to the contrary*, because it is this (together with the uncontroversial claim that *Fountain* is a work of conceptualism) that yields a solution to the meta-problem. Here is Lopes on the meta-problem and his solution to it:

Since the hard cases [like *Fountain*] do grip us, the buck passing theory of art is not a live option unless it enables us to contend with them in a serious manner. To do this, it [...] must represent the hard cases as posing a puzzle that calls upon philosophy for a response. [...] According to the buck passing theory, what makes an item a work of art is a matter of its meeting two conditions: it belongs to a kind and that kind numbers among the arts. An item is a work of art if and only if it is a product of some art kind. The question of what makes *Fountain* a work

---

<sup>111</sup> For several considerations in favour of understanding conceptualism as a genre, see, for example, Dodd [2016: 256-60; 2017: 639-40].

of art (if it is one) is replaced with a different question, of what makes it a sculpture . . . or a member of some other art kind (if it is one). The new question leaves room for puzzlement. If *Fountain* is a work of art then maybe it is a sculpture. But then it is puzzling what makes things sculptures . . . so perhaps it belongs to some other art? Which one? [Lopes 2014: 187]

On his solution, then, *Fountain* poses theoretical puzzlement (puzzlement about whether it is art) because it is puzzling which art, if any, *Fountain* belongs to. So, Lopes' view yields a unified solution to the theoretical problem and its meta-problem. (Whether this unified solution is a good solution is not my concern here.)

### 2.3. *Two problems from art history*

#### 2.3.1. *The problem of the change in the practice*

After Duchamp's submission of *Fountain* to the annual exhibition of the Society of Independent Artists, the practice of treating things as art radically changed. Not only is *Fountain* now treated as art within the practice of treating things as art—for example, it is exhibited at the Philadelphia Museum of Art, Tate Modern, and other leading art museums—but so are, for instance, turn-of-the-twenty-first-century works such as Tracey Emin's *My Bed* (her actual unmade bed and other everyday objects, such as used condoms and underwear, she had accumulated around it during a period of depression) and Martin Creed's *Work No. 227: The lights going on and off* (an empty room whose lights are on for five seconds, after which they are turned off for five seconds, to then be turned back on for another five seconds, and so on *ad infinitum*). This sample of the change is, of course, wildly under-representative of it, but it is representative enough for us to understand the problem that concerns us here: when and how did this change in the practice of treating things as art *exactly* occur? Unsurprisingly, given my argument in Section 2.1, this should be common ground among solutions to the problem of the change in the practice: Duchamp's discovery in the 1910s of what I have called Duchamp's

paradox played a key role in triggering the change in the practice.<sup>112</sup> (This is obviously consistent with views, like Beardsley's, on which *Fountain* is not art, for something's being treated as art is consistent with its not being art.) In what follows, I show, drawing on work in the social history of art by Patricia Mainardi [1993] and Thierry De Duve [2013b, 2014a, 2014b], that something more surprising should also be common ground among solutions to the problem: the conditions for someone to discover Duchamp's paradox—and thus potentially trigger the change in the practice—were in place, historically speaking, *before Duchamp was even born*. (He was born in 1887.) In doing so, I trace the *origins* of the change in the practice.

As we saw in Section 2.1, the Society of Independent Artists had a 'governing principle', which, I suggested, should be understood as follows: any member of the Society of Independent Artists had the right to exhibit their submitted work at the annual show without interference—a demand-right against the Society's board of directors. As we also saw, this governing principle was the basis on which premise (2) of Duchamp's paradox rested, for it was based on it that Mutt, as a member of the Society, was able to put the board under the obligation of exhibiting his submitted work (on pain of violating his right, and thus the governing principle). The Society did not come up with that governing principle, however, as the announcement of its first annual exhibition makes clear:

The program of the Society of Independent Artists, which is practically self-explanatory, has been taken over from the Société des Artistes Indépendants of Paris. The latter Society, whose salon is the oldest in France, has done more for the advance of French art than any other institution of its period. A considerable number of the most prominent artists of the present generation and the preceding one established their reputation at its annual exhibitions [...]. The reason for this success is to be found in the principle adopted at its founding in 1884 and never changed: 'No jury, no prizes.' [Glackens et al. 1917, fol. 3]

---

<sup>112</sup> Note that Duchamp probably discovered the paradox before 1917, since some readymades are older than *Fountain* (e.g., *Bottle Rack* dates from 1914).

That is to say, the governing principle of the *American* Society of Independent Artists was also the governing principle of the *French* Society of Independent Artists, and the former was modelled on the latter. Given the role that the governing principle plays in Duchamp's paradox, this is an important clue as to the origins of the change in the practice, leading us to nineteenth-century Paris. What was going on in Paris in the nineteenth century that explains why the French Society of Independent Artists adopted that governing principle in 1884, when it was founded by Georges Seurat, Paul Signac and other avant-garde artists of the time? In a nutshell, the *Salon*, a state-run contemporary art exhibition that became the Western world's largest and most important exhibition in the nineteenth century.

The history of the Salon cannot be disassociated from the history of the Academy of Painting and Sculpture, an institution founded by King Louis XIV in 1648 and later merged with the Academy of Music and the Academy of Architecture to form, in 1816, the Academy of Fine Arts. The Academy (as I shall, for simplicity, refer to the institution both before and after the merging) acted as the branch of the state responsible for the arts of painting and sculpture, its members being entitled to receive stipends, studios and state commissions. Before the French Revolution, in 1789, only members of the Academy could exhibit at the Salon, which was held in the Salon Carré of the Louvre Palace for many years—hence the name ‘Salon’—and whose process of selection of works by a jury was first instituted in 1748 [De Duve 2014a]. After the French Revolution, however, artists who were not members of the Academy and so independent from it—hence ‘independent’ artists—could also exhibit at the Salon. According to Mainardi [1993], this created a tension between the Academy and independent artists that would last for almost a century.

To understand the tension, it is crucial to bear in mind that, for most of the nineteenth century, exhibiting at the Salon was critical to achieve success as an artist in France. Independent artists whose works had not been deemed worthy of being exhibited at the Salon by the jury could hardly earn a living as artists, since being rejected meant receiving far fewer

commissions from the state and the church as well as from municipalities and private collectors—such was the power of the Salon jury in nineteenth-century French culture.

In the early 1830s, after many requests from independent artists that the Salon became again an annual event, rather than a biennial event (as it had been for the last two decades), King Louis Philippe sought a compromise between independent artists and the Academy. To please independent artists, he made the Salon an annual event again; to please the Academy, he appointed it as the jury of the Salon [Mainardi 1993: 27]. In part because the Academy perceived so-called ‘history painting’—exemplified by large-scale paintings like Jacques-Louis David’s *The Death of Socrates*, exhibited at the 1787 Salon—as ranking highest in the hierarchy of genres in painting, and so-called ‘genre painting’ (exemplified by smaller-scale paintings that depict ordinary people engaged in everyday activities, such as Jean-Baptiste Greuze’s *Broken Eggs*, exhibited at the 1757 Salon) as ranking low in the hierarchy, the result of the compromise was that the Academy started rejecting submissions to the Salon in large quantities.

In 1840, for example, out of 3996 works submitted to the Salon, more than half were rejected, which led some to demand, increasingly over the years, that those whose works had been rejected should have an exhibition of their own. This would only officially happen in 1863, when, despite the jury’s verdict, Emperor Napoleon III (nephew of Napoleon Bonaparte) granted the rejects the exhibition space they needed (next to the Salon proper) to show their work, in what became known as the Salon des Refusés. It was at the Salon des Refusés that Edouard Manet’s then controversial *Le Déjeuner sur l’Herbe*—a painting depicting ordinary people having a picnic (a feature of genre painting) but measuring over two meters in height and width (a feature of history painting)—was first exhibited and, standard art history has it, modernism was born.<sup>113</sup>

---

<sup>113</sup> On Manet and the rise of modernism, see Fried [1996] and Bourdieu [2017].

The relevance of the history of the Salon for the problem of the change in the practice is that, as the events described unfolded, philosophical discussions about the *rights* of independent artists were taking place in the background. As Mainardi [1993: 20] shows, one argument commonly used to challenge the authority of the Academy over the Salon consisted in appealing to the right to freedom of expression, especially as it applied to the press, and extrapolating from it a *right to exhibit*. Such an argument was used as early as 1791, the year (post-French Revolution) that the Salon became open to all artists, as is clear from this revealing statement from the Constitutional Committee of the National Assembly:

The equality of rights that forms the basis of the constitution has permitted all citizens to express their thoughts. This legal equality should permit all artists to exhibit their work. Their pictures—those are their thoughts—and their public exhibition—this is the permission to publish. The Salon of the Louvre is the press for pictures, provided only that morals and public order are respected. [Quoted in Mainardi 1993: 20]

The same argument was used time and again throughout the nineteenth century, as was the case when the sculptor David d'Angers proposed, in 1838, the following way of easing the tension between the Academy and independent artists: a system of two exhibitions, one of which would be a non-juried exhibition, on display for six months at a time, in which any artist could show two works, and the other a juried exhibition, taking place every ten years, in which the best works of the decade would be shown. With such a system in place, d'Angers argued, '[t]o exhibit would no longer be a privilege but a right' [quoted in Mainardi 1993: 28].

What I take all this to show is that the philosophical discussions about the rights of independent artists in nineteenth-century France played a key role in triggering the change in the practice. One upshot of this is that Duchamp's paradox is not as tied to the context of the American Society of Independent Artists as it may seem: it could have arisen (in a non-trivial sense) regardless of this particular context—for instance, had an independent artist in

nineteenth-century France appealed to their right to exhibit a readymade. (I shall return to this link between the Mutt affair and the origins of change in the practice in Section 4.2.)

### 2.3.2. *The problem of the two avant-gardes via Dada*

There is a clear link between the problem of the change in the practice and the problem that concerns us here. The reason is straightforward: the avant-gardes are groups of participants in the practice of treating things as art whose activities in a particular period either anticipate or bring about significant changes in that practice. In art history, it is standard to distinguish the *historical* avant-garde (to refer to the activities of certain groups of participants in the practice in the early 1910s through the mid-1920s) from the *neo*-avant-garde (to refer to the activities of certain groups of participants in the practice in the mid-1950s through the mid-1970s) and to ask what their relationship is. More specifically, it is standard to ask: what did the neo-avant-garde do that had not already been done, to a large extent, by the historical avant-garde? This is the problem of the two avant-gardes. Why is it a problem at all? In part because some of the novel ‘devices’ used by movements of the historical avant-garde—such as Dada and Constructivism<sup>114</sup>—were reused extensively by movements of the neo-avant-garde—such as Conceptualism—and this empirical fact alone raises the question of what the relationship between the two avant-gardes is. Two of the devices in question are the *readymade* and the *monochrome*.

The readymade made its first appearance, as we already saw, in the context of (New York) Dada, when Duchamp, in the 1910s, selected certain everyday objects and went on to

---

<sup>114</sup> For those unfamiliar with Constructivism, here is a brief explanation of some its basic tenets, and the context in which it emerged, by the art historian and Constructivism scholar Christina Lodder [1983: 1]: ‘[Constructivism] was a direct response to the experience of the Russian Revolution of 1917 and of the ensuing Civil War. The far-reaching and utopian aspirations which inspired those artists who adhered to Constructivism were embodied in works like [Vladimir] Tatlin’s Monument to the Third International [...]. Ever since the West first became aware of the artistic developments of post-revolutionary Russia in 1922 [...], it has viewed Russian Constructivism pre-eminently as an art movement. In reality it was something much wider: an approach to working with materials, within a certain conception of their potential as active participants in the process of social and political transformation.’

sign, date and title them.<sup>115,116</sup> Several decades later, however, it made a reappearance, most notably in the context of (proto- and proper) Conceptualism, when many members of the avant-garde of the 1950s through the 1970s—the neo-avant-garde—either reused it or extended it.<sup>117</sup> As I understand these terms here, a *reuse* of the readymade is a more or less narrow interpretation of its implications for the practice of treating things as art, while an *extension* of the readymade is a wider interpretation of its implications. The distinction is best understood by way of examples. Robert Morris' *Three Rulers* (three ordinary rulers hanging on a wall), Daniel Spoerri's *Prose Poems* (the leftovers of a meal on a wooden board, including plates, a knife and a bottle) and Joseph Kosuth's *One and Three Chairs* (a chair, a photograph of that chair, and a dictionary definition of 'chair') are reuses of the readymade. This is because of the narrow way in which they interpret its implications, along the lines of: 'If that urinal may be treated as art, so may other everyday objects that I select, such as *these* three rulers'—as in Morris' iteration of Duchamp's paradox. On the other hand, John Cage's *4'33"* (a three-movement composition, for any instrument, whose score instructs the musicians to not play their instruments), Robert Rauschenberg's *Erased de Kooning Drawing* (a piece of paper with barely any traces of a drawing by Willem de Kooning, erased as it was by Rauschenberg) and Sherrie Levine's *Shoe Sale* (a sale of seventy-five pairs of black shoes, each for two dollars, carried out in New York City) are extensions of the readymade. This is because of the wider way in which they interpret its implications: 'If that urinal may be treated as art, so may other things beyond everyday objects that I select, such as *these* four minutes and thirty-three seconds of silence'—as in Cage's iteration of Duchamp's paradox.

---

<sup>115</sup> For an account of the history of Dada written by a Dadaist, see Richter [1965]. For a history of New York Dada in particular, see Naumann [1994] and Hopkins [2016].

<sup>116</sup> There was at least one New York Dadaist who followed Duchamp's lead closely in the 1910s: the Baroness Elsa von Freytag-Loringhoven, who in 1917 (perhaps in the collaboration with another New York Dadaist, Morton Schamberg) created her own readymade—*God* (a plumbing pipe mounted on a wooden mitre box).

<sup>117</sup> For more on Conceptualism, see, for example, Lippard and Chandler [1971], Buchloh [1990], and Alberro [1999]. On the reception of Duchamp's work by the neo-avant-garde, see Buchloh et al. [1996] and Hopkins [2006].



As for the monochrome, it made its first appearance in the context of Constructivism, when Alexander Rodchenko, in the 1920s, painted the triptych *Pure Colours: Red, Yellow, Blue*, a work consisting of three canvases, each of which with its surface wholly and uniformly covered by paint of a single primary colour. Here is what Rodchenko would later say about it: '[In 1921] I reduced painting to its logical conclusion and presented three canvases: red, blue, and yellow. I affirmed: this is the end of painting. These are the primary colours. Every plane is a discrete plane and there will be no more representation.'<sup>118</sup> Like the readymade, the monochrome made a reappearance several decades later; for example, when Yves Klein, in the 1950s and early 60s, created his blue monochrome paintings.

What is then the relationship between these two avant-gardes? In the book *Theory of the Avant-Garde*, the literary critic and historian Peter Bürger [1984, 2010] proposed an influential solution to the problem: that by reusing or extending the devices that the historical avant-garde had used to carry out its (failed) project, such as the readymade and the monochrome, the project of the neo-avant-garde was derivative (and also failed). In what did the project of the historical avant-garde consist, according to Bürger? In a nutshell, it consisted in achieving the aim that united Dada, Constructivism and the other movements of the historical avant-garde such as Surrealism<sup>119</sup>: to challenge the so-called 'autonomy' of art—the (then prevalent) view that art has intrinsic value<sup>120</sup>—and thereby challenge the institutions of art that upheld it.

Bürger's solution has been heavily criticised (and rightly so, on my view), most notably by the art historian Benjamin Buchloh [1984, 1986, 2000], on the grounds that the historical avant-garde, in all its diversity, cannot be understood in terms of a single overarching aim,

---

<sup>118</sup> Quoted in Buchloh [1986: 44].

<sup>119</sup> For those unfamiliar with Surrealism, here is a brief explanation of some its basic tenets, and the context in which it emerged, by the (already mentioned) art historian and Surrealism scholar David Hopkins [2004: xiv]: 'Surrealism [...] was officially born in 1924 and had virtually become a global phenomenon by the time of its demise in the later 1940s. Committed to the view that human nature is fundamentally irrational, Surrealist artists such as Max Ernst, Salvador Dalí, Joan Miró, and André Masson conducted an often turbulent love affair with psychoanalysis, aiming to plumb the mysteries of the human mind.'

<sup>120</sup> See Gaiger [2009: 52-3].

which is precisely what Bürger proposed. The soundness of this objection is compatible, however, with Bürger's solution being the right solution to a *sub*-problem of the problem of the two avant-gardes: what did the neo-avant-garde do—through Conceptualism—that had not already been done, to a large extent, by the historical avant-garde through Dada? This is what I called in Section 1 the problem of the two avant-gardes *via Dada*, the most influential of the sub-problems making up the problem of the two avant-gardes.<sup>121</sup>

In what follows, I contrast a Bürgerian solution—that is, a solution in the spirit of Bürger's—to that sub-problem with an alternative to it, which I call a Fosterian solution, one in the spirit of the alternative solution to the problem of the two avant-gardes suggested by the art historian Hal Foster [1996].<sup>122</sup> What is it for a solution to the problem of the two avant-gardes *via Dada* to be a Bürgerian solution? It is for it to put forward at least three theses: that Dada had an overarching aim that it sought to achieve (a project), that its project failed, and that the project of the neo-avant-garde through Conceptualism was derivative of the project of Dada and also failed.

Consider the following two extracts from interviews with Duchamp in 1968 (the interviews were given to the BBC and the magazine *Art News*, respectively):

*Interviewer:* What you were also attempting to do [with the readymades], as I understand it, was to devalue the art as an object simply by saying: 'If I say it is a work of art, that makes it a work of art.'

*Duchamp:* Yeah, but the word 'work of art', you see, is not so important for me. I don't care about the word 'art' because it has been so discredited.

*Interviewer:* But you in fact contributed to the discrediting, didn't you, quite deliberately?

---

<sup>121</sup> An example of another of these sub-problems is what might be called the problem of the two avant-gardes *via Constructivism*: what did the neo-avant-garde do that had not already been done, to a large extent, by the historical avant-garde through Constructivism? For more on this sub-problem, see especially Buchloh [1986].

<sup>122</sup> Why spell out a Bürgerian solution to the problem rather than Bürger's actual solution? Because I find myself unable to reconstruct Bürger's actual solution in clear terms, since there are gaps in his statement of it and these cannot be filled based on his more recent statement of the solution [in Bürger 2010: 707-14]. That said, it is worth emphasising that the Bürgerian solution that I spell out in what follows preserves the spirit of Bürger's actual solution. (*Mutatis mutandis*, the same can be said about why I spell out a Fosterian solution rather than Foster's actual solution.)

*Duchamp*: Deliberately, yes. So, I want to get rid of it, [in the] way many people today have done away with religion [...].<sup>123</sup>

*Interviewer*: Do you think that the act of creating a readymade is an act of art?

*Duchamp*: I wouldn't say so, no. The fact that they are regarded with the same reverence as objects of art probably means that I have failed to solve the problem of trying to do away with art.<sup>124</sup>

In addition to these, consider also Duchamp's obscure remarks in a note from 1914 (wholly neglected by analytic philosophers) entitled 'Principle of contradiction', about what is today best known as the law of non-contradiction<sup>125</sup>:

From the Principle of Contradiction, defined only by these 3 words: i.e., *Counderstanding of Opposites*, abrogate all sanctions establishing the *proof\** of *this* in relation to its abstract opposite *that*. [T]hus understood, the principle of contradiction insists on the abstract uncertainty the contrast, to the concept A, of its opposite, B. develop. Again here, the principle of contradiction remains constant i.e., contrasts 2 more opposites. *By nature, it can contradict its own self* and require

[1<sup>st</sup>] either a return to a logical, noncontradictory continuation. (Plato ..... [sic]

[2<sup>nd</sup>] or the very contradiction, by the principle of contr, of the statement A., against B no longer A's opposite, but different (the no. of B's is infinite, analogous to the plans of a game which would no longer have rules.)

After having multiplied B. to infinity the result eventually no longer validates the statement of A. (A, theorem, is no longer formulated, nor formulatable.) It liberates the word from *definition*, from *ideal meaning*.<sup>126</sup>

Interestingly, in 1927, Duchamp created a lesser-known work which, either seriously or playfully, he assumed to be a counter-example to the law of non-contradiction: a door, in a corner of his small apartment, that served both as the door of the bedroom and as the door of the bathroom, so that when (say) the bedroom door was closed, the bathroom door was open—a case in which, he assumed, the door was and was not closed.<sup>127</sup>

---

<sup>123</sup> See [vimeo.com/165727042](https://vimeo.com/165727042) (17:23–18:02).

<sup>124</sup> Quoted in Camfield [1989: 97].

<sup>125</sup> *Law of non-contradiction*: For any proposition *p*, it is not the case that both *p* and not-*p* (at the same time and in the same sense).

<sup>126</sup> Quoted in Nesbit and Sawelson-Gorse [1996: 139].

<sup>127</sup> See Naumann [1990: 34-5].

In light of Duchamp's remark that he wanted to get rid of 'art' as well as his remarks about and longstanding interest in contradictions (a common interest of Dadaists<sup>128</sup>), the first step of a Bürgerian solution suggests itself: Dada aimed to *eliminate* the concept of art<sup>129</sup>, and it sought to achieve this aim by showing—through its primary device, the readymade—that the concept of art is *inconsistent*. Before fleshing this out, we need to get clear on what an inconsistent concept is supposed to be.

To draw on Kevin Scharp's account, let us assume that a concept is inconsistent iff its constitutive principles are jointly inconsistent (with a set of facts), where these principles are constitutive for that concept if they partially determine the meaning of the term which expresses the concept and the identity of the concept [Scharp 2013: 36]. Scharp's toy example of an inconsistent concept is the one expressed by (the made-up word) 'rable', which has P1 and P2 as constitutive principles:

P1: 'Rable' applies to  $x$  if  $x$  is a table.

P2: 'Rable' does not apply to  $x$  if  $x$  is a red thing.

Suppose that that we possess the concept of rable and that possessing it involves accepting both P1 and P2. Given the existence of red tables, P1 entails that any red table is a rable and P2 entails that no red table is a rable, and so P1 and P2 together entail that a red table is and is not a rable—a contradiction. So, in possessing the concept of rable, we possess an inconsistent concept.

With Scharp's account in hand, we can now flesh out the first step of the Bürgerian solution. What would it be for Dada to show, through a readymade like *Fountain*, that the

---

<sup>128</sup> See Richter [1965: 9]: 'Today, more than fifty years later, the image of Dada is still full of contradictions. This is not surprising. Dada invited, or rather defied, the world to misunderstand it, and fostered every kind of confusion. This was done from caprice and from a principle of contradiction.'

<sup>129</sup> To be clear, I am not saying that Duchamp's remark about wanting to get rid of 'art' commits him to wanting to eliminate the concept of art. All I am saying is that Duchamp's remark *when combined with the spirit of Bürger's actual view* suggests a Bürgerian solution according to which Dada aimed to eliminate the concept of art.

concept of art is inconsistent? It would be for it to show that the concept has constitutive principles that, given the existence of *Fountain*, jointly lead to a contradiction. Suppose that the concept has two constitutive principles:

P1\*: 'Art' applies to  $x$  if  $x$  is  $F$ .

P2\*: 'Art' does not apply to  $x$  if  $x$  is  $G$ .

Suppose further that P1\* entails that *Fountain* is art, and that P2\* entails that *Fountain* is not art. Then P1\* and P2\* together entail that *Fountain* is and is not art.

On the Bürgerian solution, Dada sought to show that there is something wrong with the concept of art in order to achieve its aim: to eliminate the concept of art. To eliminate the concept of art would involve having no natural language terms refer to it; it would involve, as Duchamp suggested, getting rid of 'art' (and presumably 'sculpture', 'painting' and other natural language terms that refer to sub-concepts of art). But—so the second step of the Bürgerian solution goes—the project of Dada clearly failed, since not only have we not got rid of 'art' but the primary device whereby Dada sought to achieve its aim (the readymade) is now treated as falling under, and believed by many to indeed fall under, 'art'.<sup>130</sup> Crucially—so the final step of the solution goes—by reusing or extending the readymade to achieve the aim of eliminating the concept of art, the project of the neo-avant-garde through Conceptualism was derivative of the project of Dada and also failed (for the same reasons that Dada failed).

To many, this is too pessimistic a view of the historical/neo-avant-garde relationship to be acceptable. Besides, there is a prominent alternative. As alluded to, Foster proposed an

---

<sup>130</sup> For one piece of evidence that this preserves the spirit of Bürger actual solution, see Bürger [2010: 705]: 'The paradox of the failure of the avant-gardes lies without a doubt in the musealization of their manifestations as works of art, that is, in their artistic success. The provocation that was supposed to expose the institution of art is recognized by the institution as art. The institution demonstrates its strength by embracing its attackers and assigns them a prominent place in the pantheon of great artists. Indeed, the impact of the failed avant-garde extends even further. After Duchamp, not only can the everyday artefact claim the status of an artwork but the discourse of the institution is molded by the avant-gardes to a degree that no one could have predicted.'

alternative to Bürger's solution to the problem of the two avant-gardes (in general, not just via Dada), an alternative which reflects Foster's optimistic view of the historical/neo-avant-garde relationship. One central thesis of Foster's solution is that what the neo-avant-garde did, by reusing or extending such devices as the readymade and the monochrome, was to *enact* the project of the historical avant-garde for the first time [Foster 1996: 20].<sup>131</sup> This suggests an alternative, *Fosterian* solution to the problem of the two avant-gardes *via Dada*.

On the Fosterian solution, Dada aimed not to eliminate the concept of art but to *replace* it by a new one, an aim which it sought to achieve by showing—through its primary device, the readymade—that the concept of art is inconsistent. To replace the concept of art would involve endorsing a theory of art that accommodates the readymade. *This* was the project of Dada. But the project was never in fact carried out, in part because *Fountain*—the paradigm readymade—fell into obscurity soon after the Mutt affair took place in 1917 and only much later, from around the early 1960s, did the reception of Duchamp's readymades by the neo-avant-garde properly begin.<sup>132</sup> Accordingly, what the neo-avant-garde did through Conceptualism was to enact the project of Dada for the first time. In fact, the Fosterian solution fits well with Scharp's account of the (typical) stages of a conceptual revolution:

1. *Pre-revolution*: people possess and use concept X and theory T in which X serves an explanatory role (e.g., mass and Newtonian mechanics).

---

<sup>131</sup> Since Foster's solution assumes (with Bürger's) that the historical avant-garde had a single overarching aim that it sought to achieve, and is thereby (like Bürger's) subject to Buchloh's charge, I think it is also best interpreted as a solution to the problem of the two avant-gardes *via Dada*.

<sup>132</sup> In fact, the time that it took for the reception of Duchamp's readymades to begin puzzles art historians. Here is an exchange between Buchloh and Andy Warhol about it in 1985:

*Warhol*: No. I didn't know him [Duchamp] that well; I didn't know him as well as Jasper Johns or Rauschenberg did. They knew him really well.

*Buchloh*: But you had some contact with him?

*Warhol*: Well, yeah, we saw him a lot, a little bit. He was around. I didn't know he was that famous or anything like that.

*Buchloh*: At that time, the late fifties and early sixties, he was still a relatively secret cult figure who just lived here [in New York City].

*Warhol*: Well, even all the people like Barney Newman and all those people, Jackson Pollock and Franz Kline, they were not well known.

*Buchloh*: In retrospect it sometimes seems unbelievable that the reception process of Duchamp's work should have taken so long. [Buchloh 1996: 37].

2. *Early revolution*: people discover that X is an inconsistent concept; they have some idea of which situations cause problems for those who use X; because of these problems, doubt is cast on the explanatory force of X and the acceptability of T as fundamental theory; however, without an alternative, people still use T and X.
3. *Late revolution*: new concepts (say  $Y_1, \dots, Y_n$ ) are proposed and a new theory (say U) is proposed in which the  $Y_i$ s serve an explanatory role (e.g., relativistic mass and proper mass in relativistic mechanics); U reduces to T in familiar cases, and the  $Y_i$ s agree with X on familiar cases; U is used to determine the cases in which it is acceptable to use T; at this point the conceptual repertoire and language have been extended.
4. *Post-revolution*: U has replaced T as the accepted fundamental theory, and the  $Y_i$ s have replaced X as the accepted fundamental concepts; people might or might not still use T (and thus X) in certain cases (e.g., phlogiston theory has been totally superseded, but Newtonian mechanics is still indispensable for everyday situations). [Scharp 2013: 137]

Applied to the concept of art, Stage 1 would correspond, according to the Fosterian solution, to the period before the historical avant-garde (which, incidentally, coincides with when aesthetic theories of art were prominent). Stage 2 would correspond to the period of the historical avant-garde, when Duchamp, in the context of Dada, discovered that the concept of art has constitutive principles that, given the existence of the readymade, jointly lead to a contradiction (e.g., *Fountain* is and is not art). Stage 3 would correspond to the period of the neo-avant-garde, when Conceptualists (members of Conceptualism) enacted the project of Dada for the first time (which, incidentally, coincides with when Dickie [1969], drawing on Arthur Danto [1964] and Maurice Mandelbaum [1965], developed the first ever institutional theory of art). Stage 4 would correspond to the period in which we live now.

As should be clear, whether the Fosterian and Bürgerian solutions are good solutions to the problem of the two avant-gardes via Dada is not my concern here. My concern here has only been to spell out (improved versions of) the two most prominent solutions to the problem in the literature, so that later, in Section 5, it becomes clear how my solution to it—a part of my solution to the Duchamp’s box problem—differs from those.

### 2.3.3. *We need a unified solution to these two problems*

As far as I am aware, there is no attempt in the literature to provide a unified solution to the problem of the change in the practice and the problem of the two avant-gardes via Dada. But we need a unified solution. To see why, recall that the former problem is about when and how the radical post-*Fountain* change in the practice, with its origins in nineteenth-century France, exactly occurred. At the same time, notice that the latter problem is about the relationship between the activities of Dadaists and Conceptualists, participants in the practice of treating things as art whose activities, in different periods of the twentieth century, either anticipated or brought about *the very same radical change in the practice with which the former problem is concerned*. Thus, at their core, both problems concern the same phenomenon.

### 2.4. *The practical problem and its meta-problem*

How, if at all, should *Fountain* be appreciated as art? This is the practical problem. Properly understood, the problem breaks into two questions:

Qa: May *Fountain* be appreciated as art? In other words, is it appropriate to appreciate *Fountain* as art?

Qb: If so, how should *Fountain* be appreciated as art? In other words, what is the appropriate way (or, if there is more than one, what are the appropriate ways) to appreciate *Fountain* as art?

For those who propose to solve the practical problem by answering Qa in the positive, Qb arises. Unless they answer the latter, they cannot rightly claim to have provided a solution to the problem. For those who propose to solve the problem by answering Qa in the negative, Qb does not arise, in which case they can rightly claim to have provided a solution to the problem without answering Qb.



As I see it, the current debate over the practical problem is best understood as a set of different reactions to a three-premise argument to the conclusion that *Fountain* should *not* be appreciated as art—that is, to the conclusion that the right answer to Qa is negative. According to the first two premises of the argument,

- (1) (For any  $x$ ) If  $x$  is appreciated as art, then the aesthetic properties of  $x$  are appreciated.
- (2) (For any  $x$ ) If the aesthetic properties of  $x$  are appreciated, then they are appreciated on the basis of (at least some of) the perceptual properties of  $x$ .

Together, these premises entail that to appreciate  $x$  as art involves appreciating  $x$  on the basis of the perceptual properties of  $x$ . To appreciate *Le Déjeuner sur l'Herbe* as art, for example, involves appreciating it on the basis of its perceptual properties—the properties of it available to the senses—such as the sketch-like (non-detailed) way in which some of the trees in the composition are depicted, a property that we can detect by looking at Manet's painting.

Is appreciating *Fountain* as art on the basis of its perceptual properties the way in which it should be appreciated as art? Many have remarked that it *can* be appreciated in this way (where the 'can' here is that of possibility, not permissibility). Here is Dickie making this remark (in the context of an objection of Ted Cohen's [1973] to his institutional theory which need not concern us here):

But why cannot the ordinary qualities of *Fountain*—its gleaming white surface, the depth revealed when it reflects images of surrounding objects, its pleasing oval shape—be appreciated? It has qualities similar to those of works by [Constantin] Brancusi and [Henry] Moore which many do not balk at saying they appreciate. [Dickie 1974: 42]

It seems hard to deny that it is possible to appreciate *Fountain* as art in the way that Dickie suggests. In fact, as the art historian William Camfield [1989: esp. 39-42] showed, pointing out that *Fountain* could be appreciated as art on the basis of its (allegedly) pleasing perceptual properties was one way in which the art status of *Fountain* was defended by those in Duchamp's

circle in 1917, most notably by the American artist Louise Norton [1917] in her article in *The Blind Man* (the obscure, avant-garde magazine co-edited by Duchamp in which, as mentioned in Section 1, the Mutt affair was documented). Nevertheless, many object to appreciating *Fountain* in this way. Here is Arthur Danto raising this objection, precisely after quoting Dickie's remark above:

These [i.e., the gleaming white surface, etc., to which Dickie refers] *are* qualities of the urinal in question, as they are qualities of any urinal made of white porcelain, which do resemble certain qualities of [Brancusi's] *Bird in Flight*. But the question is whether the artwork *Fountain* is indeed identical with that urinal, and hence whether those gleaming surfaces and deep reflections are indeed qualities of the artwork. [Ted] Cohen has supposed that Duchamp's work is not the urinal at all but the gesture of exhibiting it; and the gesture, if that indeed is the work, has no gleaming surfaces to speak of, and differs from what Moore and Brancusi did roughly as gestures differ from bits of brass and bronze. But certainly the work itself has properties that urinals themselves lack: it is daring, impudent, irreverent, witty, and clever. What would have provoked Duchamp to madness or murder, I should think, would be the sight of aesthetes mooning over gleaming surfaces of the porcelain object he had manhandled into exhibition space: 'How like Kilimanjaro! How like the white radiance of Eternity! How Arctically sublime!' [Danto 1981: 93-4]

What exactly is the objection to appreciating *Fountain* as art on the basis of its perceptual properties? I think that there are two ways of understanding it, one of which is stronger (more committal) than the other.

On the stronger, non-normative way of understanding the objection, *Fountain* is *not* the kind of thing that can be appreciated as art on the basis of its perceptual properties, simply because it does not have any perceptual property. So, those who take themselves to be appreciating *Fountain* as art by appreciating the perceptual properties of the urinal are not appreciating *Fountain* at all—they are simply confused. On the weaker, normative way of understanding the objection, *Fountain* is the kind of thing that *can* be appreciated as art on the basis of its perceptual properties, but it *should not* be so appreciated. So, while those who appreciate *Fountain* as art by appreciating the perceptual properties of it are not confused—

*Fountain* does have perceptual properties (viz., those it shares with the urinal)—they are doing something they should not be doing.

Since I think that it is implausible to postulate that those who appreciate *Fountain* as art by appreciating the perceptual properties of the urinal are confused, I take it that the best way of understanding the objection to appreciating *Fountain* in that way is *normative*. Given this objection,

- (3) *Fountain* should not be appreciated as art on the basis of its perceptual properties.

Together with (1) and (2), the argument to the conclusion that the right answer to Qa is negative is then as follows:

- (1) (For any  $x$ ) If  $x$  is appreciated as art, then the aesthetic properties of  $x$  are appreciated.  
(2) (For any  $x$ ) If the aesthetic properties of  $x$  are appreciated, then they are appreciated on the basis of (at least some of) the perceptual properties of  $x$ .  
(3) *Fountain* should not be appreciated as art on the basis of its perceptual properties.  
Therefore,  
(4) *Fountain* should not be appreciated as art.

There have been three main reactions to this argument, each of which corresponds to a different solution to the practical problem. One reaction is to accept all its premises and thereby its conclusion. This is Beardsley's solution to the problem (or at least, I suggest, the most charitable way of interpreting it). On his solution though, while *Fountain* should not be appreciated as art, it may be appreciated as art *theory*.<sup>133</sup> Another reaction is to accept premises (1) and (2) but reject premise (3). This is the solution to the problem held, for instance, by the art critic Clement Greenberg. On this solution, there is no reason to doubt that *Fountain* may

---

<sup>133</sup> See Beardsley [1981: esp. 29-34; 1983: esp. 60].

be appreciated as art, since *Fountain* may be appreciated as art on the basis of its perceptual properties.<sup>134</sup> However, neither of these is the dominant solution to the practical problem.

According to the dominant solution, the right reaction to the argument is to accept premise (3) but reject either premise (1) or (2). On this solution, there is also no reason to doubt that *Fountain* may be appreciated as art—to doubt, that is, that the right answer to Qa is positive. So, Qb arises: how should *Fountain* be appreciated as art? By being appreciated (merely) on the basis of its *non-perceptual* properties, such as its daring, impudence and wit. This is the solution held most notably by Danto [1981] and James Shelley [2003]. However, its proponents disagree sharply among themselves about how to best interpret the view, since some reject the conjunction of premises (1)–(2) by rejecting only (1), others do so by rejecting only (2), and still others reject the conjunction by rejecting both (1) and (2).

Danto rejects premise (1) but accepts premise (2).<sup>135</sup> This yields a version of the view according to which, when one appreciates *Fountain* (as art) as one should—namely, by appreciating (merely) its non-perceptual properties—one is appreciating it *non-aesthetically*. Why? Because, by (2), if one appreciates *Fountain* by appreciating (merely) its non-perceptual properties, then one is appreciating the *non-aesthetic* properties of it. On the other hand, Shelley rejects premise (2) but accepts premise (1). This yields another version of the view, one according to which, when one appreciates *Fountain* (as art) as one should, one is appreciating it *aesthetically*. Why? Because, by (1), if one appreciates *Fountain* as art, then one is appreciating the aesthetic properties of it, in which case *Fountain*'s daring, impudence and wit are *aesthetic* properties of it (albeit non-perceptual ones). Moreover, there are proponents of the view, like Noël Carroll [2004], who reject both premises (1) and (2), leaving it compatible with either version.<sup>136</sup>

---

<sup>134</sup> See Greenberg [1986].

<sup>135</sup> For another prominent proponent of this version of the view, see Binkley [1970].

<sup>136</sup> More recently, Costello [2013] has defended another solution to the practical problem, namely, a view according to which the right reaction to the argument is (like Shelley's version of the dominant solution) to accept premise (1) but reject premise (2), *and* (unlike any version of the dominant solution) reject premise (3). More

Given the parallel between the discussion in this section and the discussion in Section 2.2 (on the theoretical problem and its meta-problem), it should come as no surprise that any solution to the practical problem faces the challenge of explaining why *Fountain* poses the practical problem in the first place—what I have called the practical meta-problem. In fact, just like we should expect for there to be a unified solution to the theoretical problem and its meta-problem, we should expect for there to be a unified solution to the practical problem and its meta-problem: if we had a solution to the practical problem, we would expect it to yield, or at least provide the basis for, a solution to the meta-problem; if we had a solution to the meta-problem, we would expect it to yield, or at least provide the basis for, a solution to the practical problem.

It is not standard in the literature, however, to attempt to provide a unified solution to the practical problem and its meta-problem (in part because it is not standard to distinguish the two problems). Nevertheless, there are solutions to the practical meta-problem in the literature. Moreover, some solutions to the practical problem yield a solution to the practical meta-problem along the lines of one already in the literature, thus providing a unified solution to both problems—which is exactly what we need. Let me elaborate on both of these points.

One of the solutions to the meta-problem in the literature is Robert Hopkins' [2007].<sup>137</sup> Hopkins' [2007: esp. 62] explanation for why *Fountain* poses practical puzzlement—puzzlement about how, if at all, it should be appreciated as art—is as follows. *Fountain* is the kind of thing that *can* be appreciated as art on the basis of its perceptual properties (where, again, the 'can' here is that of possibility, not permissibility). Because of this, and the contexts in which we find *Fountain* (e.g., art museums), we create an expectation: that, when we

---

specifically, Costello argues, with Shelley, that when one appreciates *Fountain* (as art) as one should, one is appreciating it aesthetically. However, contra Shelley and all other proponents of the dominant solution, he argues that it is not case that one should appreciate *Fountain* as art *merely* on the basis of its non-perceptual properties, since he thinks that *Fountain* should be appreciated as art *both* on the basis of its non-perceptual properties (e.g., its wit) *and* on the basis of its perceptual properties (e.g., its gleaming surface).

<sup>137</sup> See Lopes [2007] for another.

encounter *Fountain*, it will be satisfying for us to appreciate it as art through the senses (i.e., on the basis of its perceptual properties). But *Fountain* frustrates this expectation, since there is nothing satisfying in appreciating it as art through the senses—hence the practical puzzlement posed by *Fountain*.

Some solutions to the practical problem yield a solution to the meta-problem along the lines of Hopkins'. Consider solutions that accept premise (3), such as the dominant solution and Beardsley's solution. On these solutions, as we saw, while *Fountain* can be appreciated as art on the basis of its perceptual properties, it should not be so appreciated. So, proponents of these views can show, by appealing to the basic tenets of their solutions (and the claim that we find *Fountain* in the relevant contexts), that we create the relevant expectation and that *Fountain* frustrates it. More specifically, they can show that we create the expectation because, given their solutions, *Fountain* is the kind of thing that can be appreciated as art in the relevant way (and we find it in the relevant contexts), and they can show that *Fountain* frustrates the expectation because, given their solutions, to appreciate *Fountain* as art in that way is to do something that we should not do. Thus, these views yield a unified solution to the practical problem and its meta-problem. (As before, whether this unified solution is a good solution is not my concern here.)

## 2.5. *Why do we need a solution to the Duchamp's box problem?*

A few conclusions emerge from the preceding discussion. For our purposes here, three of them are especially relevant, but let us first concentrate on only two. One conclusion from Section 2.2 is that we need a unified solution to the theoretical problem and its meta-problem, given the undeniable link between them. One conclusion from Section 2.4 is that we need a unified solution to the practical problem and its meta-problem, given the undeniable link between them. So far, so good. But why think that we need a unified solution to all these four

problems? Because we need a unified solution to Duchamp's paradox and the theoretical and practical problems. To see this, consider Duchamp's paradox again:

*Duchamp's paradox*

- (1) (For any  $x$ ) One must not: treat  $x$  as art if  $x$  is not art. (Do Not TAA)
- (2) *Fountain* may be treated as art (i.e., it is permissible to treat it as art).

Therefore,

- (C) *Fountain* is art.

Furthermore, notice that this is true: if one appreciates  $x$  as art, then one is treating  $x$  as art. So, if one appreciates  $x$  as art, but  $x$  is not art, then Do Not TAA is violated. This implies that, if Do Not TAA holds, there is a norm governing appreciation within the practice of treating things as art: do not appreciate as art that which is not art. More formally,

*Do Not AAA*: (For any  $x$ ) One must not: appreciate  $x$  as art if  $x$  is not art.

The upshot of this is twofold. If Do Not AAA holds and *Fountain* is not art, then *Fountain* should not be appreciated as art—in which case the *practical* problem would be solved. If Do Not AAA holds and *Fountain* may be appreciated as art, then *Fountain* is art—in which case the *theoretical* problem would be solved. Duchamp's paradox thus reveals that there is an intriguing link between itself, the theoretical problem, *and* the practical problem. Accordingly, any solution that can only take on some of these problems is a solution that misses the forest for the trees. I conclude that we need a unified solution to Duchamp's paradox and those two problems, and so a unified solution to these three problems and the meta-problems.

Let us now turn to the other relevant conclusion that emerges from the preceding discussion. One conclusion from Section 2.3 is that we need a unified solution to the problem of the change in the practice and the problem of the two avant-gardes via Dada. After all, they

both concern the very same radical change in the practice. But why think that we need a unified solution to these two problems and the other five problems? Because, as I argued in Section 2.3.1 drawing on work in the social history of art, there is an intriguing link between Duchamp's paradox and the origins of the change in the practice—a link which suggests that there was something that Duchamp discovered in the 1910s, through his paradox, which provides the basis for an account of the mechanisms whereby the change in the practice occurred. Any unified solution to those two problems that ignores this link is thus a solution that fails to get to the heart of the matter. This is why we need a unified solution to those two problems and the other five problems, and so a solution to the Duchamp's box problem.

### 3. THE NORMATIVE INDETERMINACY HYPOTHESIS

The Duchamp's box problem is the problem of finding a unified solution to the seven problems that make it up. In Section 2, we saw that Duchamp's paradox is the problem that links up all the other problems. For this reason, Duchamp's paradox is arguably the best place to start looking for a solution to the Duchamp's box problem. In this section, I revisit Duchamp's paradox and introduce the basic framework for a non-standard solution to it (the paradox). I call the framework the *normative indeterminacy hypothesis*. In Section 4, I flesh out the hypothesis. In building a framework for a solution to Duchamp's *paradox* in this section and the next one, my aim is to lay the groundwork for Section 5, where I argue that the normative indeterminacy hypothesis yields a solution to the Duchamp's *box problem*.

#### 3.1. *Duchamp's paradox, revisited*

Consider again Duchamp's paradox, now stated—to ease the exposition to follow—as a set of jointly inconsistent propositions (rather than an argument to the conclusion that *Fountain* is art):



*Do Not TAA*: (For any  $x$ ) One must not: treat  $x$  as art if  $x$  is not art.

*F-May-be-TAA*: *Fountain* may be treated as art.

$F \sim \text{Art}$ : *Fountain* is not art.

Call this the ‘paradox set’. One type of solution to the paradox consists in (1a) showing that (at least) one of the propositions in paradox set is false, and in (1b) providing a convincing explanation for why that proposition appears, or appeared, to be true when considered on its own (i.e., independently of the paradox). Call this a ‘Type A’ solution. Another type of solution to the paradox consists in (2a) showing that the propositions in the paradox set are not jointly inconsistent, and in (2b) providing a convincing explanation for why they appear to be jointly inconsistent. Call this a ‘Type B’ solution. (For each of these types of solution, (#b) is indispensable because without it we cannot be confident that what is claimed to be going on—that one of the propositions is false, or that the propositions are not jointly inconsistent—is in fact what is going on.) To use Stephen Schiffer’s [2003] terminology, call any instance of either of these types of solution to a paradox a *happy-face solution*.

Does Duchamp’s paradox have a happy-face solution of Type B? No. Why? Because *Do Not TAA* and *F-May-be-TAA* entail, by *modus tollens*, that  $F \sim \text{Art}$  is false; *Do Not TAA* and  $F \sim \text{Art}$  entail, by *modus ponens*, that *F-May-be-TAA* is false; and, finally,  $F \sim \text{Art}$  and *F-May-be-TAA* entail that *Do Not TAA* is false. So, the propositions are indeed jointly inconsistent. For this reason, I set this type of happy-face solution aside in what follows. This leaves us with the happy-face solutions of Type A.

What is the correct happy-face solution of Type A to Duchamp’s paradox? According to the dominant solution to the paradox, that  $F \sim \text{Art}$  is false. (That this is the dominant solution to the paradox can easily be inferred from the literature on the theoretical problem.) For example, Lopes’ version of the buck-passing theory of art has this implication. On the buck-passing theory, recall,  $x$  is a work of art iff, for some kind  $K$ ,  $K$  is an art and  $x$  is a work of  $K$ . On Lopes’ version of this theory, *Fountain* belongs to a kind which, despite appearances to the

contrary, is an art—conceptualism. Given his view, not only is  $F \sim \text{Art}$  false but there is an explanation for why it appeared (say, in the 1910s), and why it still appears to some today, to be true: because *Fountain* did not, and still does not, appear to belong to any of the arts.

### 3.2. *Does Duchamp's paradox have a happy-face solution?*

The default view is that Duchamp's paradox has a happy-face solution, and more specifically a happy-face solution of Type A: that (at least) one of the propositions in the paradox set is false and that there is a convincing explanation for why it appeared (and still appears to some) to be true. My view, however, is that the paradox does not have a happy-face solution of Type A, and so that it does not have a happy-face solution at all (given that it does not have happy-face solution of Type B either). But why would the paradox not have a happy-face solution of Type A? Because none of the propositions in the paradox set is (*determinately*) false—or so I want to suggest. The challenge for my view then consists in providing a convincing explanation for why none of the propositions in the paradox set is (*determinately*) false. If I can meet this challenge, then I will have solved the paradox by providing (what Schiffer calls) an *unhappy*-face solution to it—that is, by providing an explanation for why Duchamp's paradox does not have a happy-face solution. In an important sense, the normative indeterminacy hypothesis is just such an explanation. Let me elaborate on this last point by introducing the bare-bones of the hypothesis and then returning to the paradox.

Recall that in Section 2.1 I appealed to Do Not TAA (one of the propositions in the paradox set) to make sense of the Mutt affair. There are independent reasons, however, to think that Do Not TAA governs the practice of treating things as art. Consider the following case:

*Glasses*: Suppose that, while visiting one of the rooms dedicated to avant-garde works of the 1960s at Tate Modern's permanent collection, I accidentally drop my glasses, without myself

or anyone else around me noticing it, and proceed to the next room without them. As other museum-goers enter the room I was in, they start treating my glasses as art—say, by taking photographs of them as they lie on the floor, ascribing meaning to them (in order to retrieve the artist’s intention), and so on.<sup>138</sup>

I take it that the museum-goers did something that they ought not to have done (where the relevant sense of ‘ought’ here is the fact-relative sense<sup>139</sup>)—namely, to treat my glasses as art—and that this explains why: because there is a norm governing the practice according to which we ought not to treat as art that which is not art, namely, Do Not TAA.

One might object that the museum-goers did not do anything wrong because they were not blameworthy for treating my glasses as art. After all, avant-garde works of the 1960s, the objection goes, are prone to generate such confusion. However, this objection fails to distinguish deontic facts (facts about what is required or permissible to do) from hypological facts (facts about what is blameworthy or praiseworthy).<sup>140</sup> Once we draw this distinction, Glasses stands as evidence for Do Not TAA.

I do not think, however, that Do Not TAA is the only norm governing the practice of treating things as art. Consider the following case:

*Trash*: Suppose that, after Tate Modern closes to the public for the day, a couple of museum employees are assigned the task of cleaning the rooms dedicated to avant-garde works of the 1960s. Mistaking one of the works of art for trash, they throw it into the bin.<sup>141</sup>

---

<sup>138</sup> This case is inspired by a real-life case. In 2016, Kevin Nguyen and T. J. Khayatan, two teenagers on a visit to the San Francisco Museum of Modern Art (in California), deliberately placed a pair of eyeglasses on the floor of the museum’s exhibition spaces with the aim of testing whether the museum-goers would treat them as art. The upshot was that they did treat the eyeglasses as art, as documented on Twitter through photographs of the museum-goers taken by Nguyen and Khayatan.

<sup>139</sup> For a distinction between the fact-relative sense and other senses of ‘ought’ (e.g., the evidence-relative sense), see Parfit [2011: 151-64].

<sup>140</sup> See, for instance, Zimmerman [2002: 554] and Lord [2017: 1136] for the distinction.

<sup>141</sup> This case is inspired by countless real-life cases.

I take it that the museum employees did something that they ought not to have done—namely, to treat that work of art as not art—and that this explains why: because there is a norm governing the practice according to which we ought to treat as art that which is art (equivalently: we ought not to treat as not art—say, as merely trash—that which is art). In other words, I think that this norm also governs the practice:

*Do TAA*: (For any  $x$ ) One must: treat  $x$  as art if  $x$  is art.

Together, Do TAA and Do Not TAA entail what I call the *TAA Norms*:

*TAA Norms*

One must: treat  $x$  as art iff  $x$  is art.

One may: treat  $x$  as art iff  $x$  is art.

Suppose then, by hypothesis, that the TAA Norms do govern the practice of treating things as art—a *prima facie* plausible claim. Here is a further, though controversial, claim: it is indeterminate whether *Fountain* is art.<sup>142</sup> Together, these two claims entail that it is indeterminate whether one ought to treat *Fountain* as art, and indeterminate whether one ought not to treat *Fountain* as art. In other words, given the two claims, *Fountain* gives rise to *normative indeterminacy*. To a first (rough) approximation, it is the conjunction of these two claims that I call the normative indeterminacy hypothesis.

To see how the normative indeterminacy hypothesis bears on Duchamp's paradox, recall the propositions in the paradox set:

*Do Not TAA*: (For any  $x$ ) One must not: treat  $x$  as art if  $x$  is not art.

---

<sup>142</sup> The claim is controversial but far from outlandish, since *Fountain* is regarded by many as a hard case and, as Crispin Wright points out in his work on indeterminacy, it is an 'absolutely basic datum that in general borderline cases come across as hard cases' [Wright 2001: 70].

*F-May-be-TAA*: *Fountain* may be treated as art.

*F~Art*: *Fountain* is not art.

If the normative indeterminacy hypothesis holds, it follows that the paradox has *no* happy-face solution. Why? Because, given the normative indeterminacy hypothesis, *F-May-be-TAA* and *F~Art* are both indeterminate, while Do Not TAA is (determinately) true. Thus, *none* of the propositions in the paradox set is (determinately) false, contrary to what a happy-face solution of Type A assumes (viz., that at least one of the propositions is false).

#### 4. SHARPENING THE NORMATIVE INDETERMINACY HYPOTHESIS

In Section 3, I introduced the basic framework for an unhappy-face solution to Duchamp's paradox—the normative indeterminacy hypothesis. According to it, it is indeterminate whether one ought to treat *Fountain* as art, and indeterminate whether one ought not to treat *Fountain* as art. This is an implication of the TAA Norms in conjunction with the claim that it is indeterminate whether *Fountain* is art. There are several ways one might go about fleshing out the normative indeterminacy hypothesis, however. In particular, there are several ways one might go about fleshing out the side of the hypothesis according to which it is indeterminate whether *Fountain* is art. In this section, I propose one way of fleshing it out.

Before introducing the proposal, though, a preliminary remark is in order. I shall understand indeterminacy in broadly supervenientist terms (where indeterminacy, as discussed in Chapter 2, is a more general phenomenon than vagueness, and so the indeterminacy to which *Fountain*, by hypothesis, gives rise might well have nothing to do with vagueness—I want to remain neutral about whether it does).<sup>143</sup> Accordingly, I take the claim that it is indeterminate whether *Fountain* is art to amount to the claim that there is at least one

---

<sup>143</sup> One reason to think that it has nothing to do with vagueness is that 'is art' does not seem to be sorites-susceptible (even though closely related predicates, such as 'is a painting', do seem to be).

sharpening of ‘is art’ on which *Fountain* is art and at least one sharpening of ‘is art’ on which *Fountain* is not art. Call these sharpenings, for now, the ‘*F*-is-art sharpening’ and the ‘*F*-is-not-art sharpening’, respectively.

With that in mind, here is what I do in what follows. In Section 4.1, I introduce Ásta’s [2018] conferralist account of social properties. In Section 4.2, drawing on her account, I shed new light on the Mutt affair and origins of the change in the practice (discussed in Sections 2.1 and 2.3.1, respectively). In Section 4.3, I introduce my proposal for fleshing out the normative indeterminacy hypothesis by specifying the content of both the *F*-is-art sharpening and the *F*-is-not-art sharpening.

#### 4.1. *Ásta’s conferralist account of social properties*

In developing her account of social properties, Ásta is primarily concerned with providing an account of the social properties of individuals, such as a person’s being cool, being elected president, being married, being a Real Madrid fan, and so on. However, to reveal the underlying structure of the account, it is helpful to begin introducing it by considering how it would apply to other social properties, such as football (soccer) properties.

Suppose that a certain course of action in a football game has the property of being a goal. On Ásta’s account, there are two layers of properties at play here: the property that gets conferred onto the relevant course of action and the property that whoever is doing the conferral is attempting to track. The property that gets conferred onto the course of action—the *conferred property*—is the property of being a goal, and it is the referee of the game who confers it. The property that the referee is attempting to track in the conferral—the *base property*—is a property that the ball may or may not have, namely, having crossed the goal line between the goalposts and under the crossbar (without the players of the attacking team having violated the rules of the game in that course of action). Given this, here are two possible scenarios:

Scenario 1: The ball has the relevant base property—that is, it *did* cross the goal line between the goalposts and under the crossbar—and the referee judges that it has the base property.

Scenario 2: The ball does not have the relevant base property—it did *not* cross the goal line between the goalposts and under the crossbar—but the referee judges that it has the base property.

In Scenario 1, the referee's judgment successfully tracks the presence of the base property. In Scenario 2, it does not—the referee made a mistake. On Ásta's account, however, this mistake is metaphysically irrelevant: in both scenarios, the course of action *is* a goal. Why? Because even though the referee made a mistake about the trajectory of the ball (it did not cross the goal line), the referee did not make a mistake about whether it is a goal, since it is the referee's judgment that the ball crossed the goal line, whether accurate or not, that confers the property of being a goal onto the course of action. Thus, so long as the referee judges that the ball has the base property, the course of action is a goal—even if the ball does not in fact have the base property.

Given a *conferralist* account of football properties, there are then five questions to ask to uncover the profile of any given football property:

What is the **conferred property**? *In our example, being a goal, conferred onto a course of action*

**Who** are the subjects doing the conferral? *The referee*

**What** is the attitude, state, or action of the subjects that matters? *The referee's judgment that the ball has the base property*

**When** (under what conditions) does the conferral take place? *In the context of a football game*

What is the **base property**? *Having crossed the goal line between the goalposts and under the crossbar, which the ball may or may not possess*

Turning now to the social properties of individuals, we can ask the same questions to uncover the profile of such properties as a person's being cool and being elected president. In fact, though, Ásta [2018: 16-22] thinks that being cool and being elected president are paradigmatic examples of distinct kinds of social property: the former is a paradigmatic example of (what she calls) a *communal property* while the latter is a paradigmatic example of an *institutional property*. Accordingly, she proposes a framework for uncovering the profile of any given communal property and another framework for uncovering the profile of any given institutional property. Let us consider these frameworks in reverse order.

Here is Ásta's [2018: 21] framework for institutional properties:

Conferred property: *P*

Who: *A person, entity, or group in authority*

What: *Their explicit conferral by means of a speech act or other public act*

When: *Under the appropriate circumstances (in the presence of witnesses, at a particular place, etc.)*

Base property: *The property (or properties) that those in authority are attempting to track in the conferral*

Given this framework, we can uncover the profile of the institutional property of being elected president—say, of the United States—as follows:

Conferred property: *Being elected president of the United States, conferred onto one of the candidates*

Who: *The current US vice-president, on behalf of the US Senate*

What: *The US Senate's declaration that a candidate has the base property*

When: *On January 6, following a November election, starting at 1 p.m.*

Base property: *Having received at least 270 electoral college votes for US president, which the candidate may or may not possess*

To flesh this out, suppose that a candidate has just acquired the property of being elected US president. Why did *this* candidate, rather than any other, acquire this property? Because an entity in authority, the US Senate, conferred the relevant property onto this candidate by means



of a speech act—namely, by declaring that this candidate has received at least 270 electoral college votes—in the appropriate institutional context.

As in the football properties example, the candidate may well not have the relevant base property, but what matters on the conferralist account is that those in authority perceive the candidate to have it. Why is it that what matters on the conferralist account is the *perception* that the base property is present, rather than the actual presence of it? Because the mere perception that the base property is present imposes *constraints* on and *enablements* to the candidate's behaviour and, on the conferralist account, to have *any* social property (or status)—such as the institutional property of being elected US president—just is to have relevant constraints and enablements [Ásta 2018: 29-30]. In this case, the relevant enablements are what the candidate is able to do in virtue of being perceived to have the base property—for instance, living in the White House; while the relevant constraints are what the candidate is not able to do, or what the candidate has pressure to do, in virtue of being perceived to have the base property—for instance, the candidate is not able to serve more than two terms as US president. These are *institutional* constraints and enablements.

Consider now Ásta's [2018: 22] framework for communal properties:

Conferred property: *P*

Who: *A person, entity, or group with standing*

What: *Their conferral, explicit or not, by means of attitudes and behaviour*

When: *In a particular context*

Base property: *The property (or properties) that those with standing are, consciously or not, attempting to track in the conferral*

Given this framework, we can uncover the profile of the communal property of being cool—as conferred, say, onto Ricardo, a student at my former secondary school—as follows:

Conferred property: *Being cool, conferred onto Ricardo*

Who: *A group of students at Miguel's former secondary school*

What: *The group's coordinated judgments that Ricardo has the base property*

When: *At Miguel's former secondary school (but not at the other secondary school in the same neighbourhood)*

Base property: *Being a good football player, which Ricardo may or may not have*

To unpack this, suppose that Ricardo is a new student at my former secondary school and has just acquired the property of being cool. Why did he acquire this property? Not because that group of students has the *authority* to confer him the property. Rather, he acquired it because the group has the *standing* to confer it. For any communal property, a story must be told as to why a certain person, entity, or group has the standing to confer the relevant communal property onto someone. In the case of Ricardo, the story might go like this. That group of students has the standing to confer the property of being cool because all members of the group share the property of being popular (at my former secondary school)—itself a communal property whose profile can be uncovered using the framework for communal properties (e.g., perhaps each member of the group is popular because each is perceived, by all other students at school, to have the base property of being funny). As a result, they have the power, in that context, to confer the property of being cool onto Ricardo.

As with institutional properties, being conferred the communal property of being cool—in virtue of being perceived by those with standing to have the relevant base property—imposes constraints on and enablements to that person's behaviour. Ricardo's being cool enables him to do things that the other, non-cool students cannot do, such as jumping the queue at the school's cafeteria every time he asks the students ahead of him if he can jump it. But being cool also constrains him—for instance, there is pressure for him to always play football at school whenever the other students are playing, even though most of the time he does not want to play. These are *non-institutional, communal* constraints and enablements.

To close, it is worth noting that not all social properties are like being US president, which is *merely* an institutional property, and being cool, which is *merely* a communal property. For some social properties, P, there is both the institutional property of being P and the communal property of being P. In her application of the conferralist account to gender, for instance, Ásta argues that there is both the institutional property of being a woman (man) and the communal property of being a woman (man).

#### 4.2. *The Mutt affair and the origins of the change in the practice, revisited*

Ásta's conferralist account allows us to shed new light on the Mutt affair and the origins of the change in the practice of treating things as art (discussed in Sections 2.1 and 2.3.1, respectively). To begin understanding why, recall that in Section 2.3.1, when tracing the origins of the change in the practice, we saw that the seeds for the change were planted, so to speak, before Duchamp was even born (with the result that, historically speaking, Duchamp's paradox could have been discovered in the nineteenth century by someone other than Duchamp). Who were they planted by? By the nineteenth-century artists who, in their struggle to get their works accepted for exhibition at the Salon, fought a political battle against the Academy of Fine Arts for the control over the Salon—the so-called *independent* artists. Given Ásta's conferralist account, we can make sense of at least one source of the tension between independent artists and the Academy, and then link this tension to the Mutt affair.

As we saw in Section 2.3.1, independent artists are so called because they were not members of the Academy, an institution that enjoyed full control over the Salon until the early 1790s, when (post-French revolution) independent artists were given permission by the state to submit their work to the Salon for jury review. As I see it, members of the Academy—known as academicians—had a social property that independent artists lacked: the *institutional property of being artists*. Why? Because an entity in authority, the Academy, had conferred the relevant property onto the former, but not the latter, individuals by means of the public act of

appointing them as members of the Academy. Here is the profile of the institutional property of being an artist in this context:

Conferred property: *Being an artist, conferred onto some individual*

Who: *The Academy of Fine Arts*

What: *The Academy's declaration that the individual has the base property*

When: *In the appropriate context in Paris in the nineteenth century*

Base property: *Being a member of the Academy*

On the other hand, independent artists (merely) had the *communal property of being artists*, which, I suggest, the various members of the Salon jury—a group with standing in that context—had conferred onto them by reviewing their submitted work and judging that they had the relevant base property, namely, I suggest, the property of self-identifying as artists. To put it another way, here is the profile of the communal property of being an artist in this context:

Conferred property: *Being an artist, conferred onto some individual*

Who: *The members of the Salon jury, collectively*

What: *The members' review of the work submitted by the individual, and their coordinated judgments that the individual has the base property*

When: *In the context of the Salon in nineteenth-century Paris*

Base property: *Self-identifying as an artist*

The fact that academicians had the institutional property of being artists, while independent artists merely had the communal property, was a source of the tension because possessing the former status consisted in having enablements that those who merely had the latter status did not have. For instance, at times in the history of the Salon in the nineteenth century, academicians were able to exhibit at the Salon without their work being reviewed by the jury—a privilege that was oppressive to independent artists.

Recall now that in Section 2.1, when analysing the Mutt affair, we saw that, according to the by-laws of the American Society of Independent Artists, anyone could become a member of it—all that one needed to do was to pay the required six dollars—and that anyone who became a member would thereby have the right to exhibit their submitted work at the annual show, without there being any jury to review it. Given the plausibility of my analysis of the tension between academicians/independent artists, we can shed new light on the close historical link between that tension and the Mutt affair as follows: acquiring the property of being a member of the American Society of Independent Artists amounted to acquiring, in that context, the institutional property of being an artist, which consisted in having the relevant constraints and especially enablements, the most notable of which was the right to exhibit. Accordingly, we can uncover the profile of the institutional property of being an artist in this context thus:

Conferred property: *Being an artist, conferred onto some individual*

Who: *The American Society of Independent Artists*

What: *The Society's declaration that the individual has the base property*

When: *In the context of the Society's annual exhibition in 1917*

Base property: *Being a member of the Society*

Why is this significant? Because, as I see it, a person's acquiring the institutional property of being an artist in this context—in virtue of being perceived by the Society to have the base property—grants that person the *authority* to confer the *institutional property of being art* onto their work, by means of the public act of submitting it to the Society's annual exhibition by the required date in 1917. If this is right, the institutional property of being art in this context has the following profile (I leave discussion of the base property for Section 4.3):

Conferred property: *Being art, conferred onto some work*

Who: *A member of the American Society of Independent Artists, and so someone with the institutional property of being an artist*

What: *That member's public act of submitting the work to the Society's annual exhibition*

When: *In the context of the Society's annual exhibition in 1917*

Base property: *P*

Like any other social property on the conferralist account, for something to have the social property of being art just is for it to have the relevant constraints and enablements. More specifically in this case: for something to have the social property of being art just is for it to impose constraints on and enablements to the behaviour of those who—as participants in the practice of treating things as art—engage with it. For our purposes here, one of these enablements is particularly noteworthy in light of the Mutt affair: if a member of the Society submitted their work to the annual exhibition, then it would be *permissible* for those who engaged with it *to treat it as art*—for instance, it would be permissible for Society's board of directors to exhibit it (and indeed, in their specific case, it would be required for them to exhibit it, on pain of violating the member's right to exhibit).

#### 4.3. *The F-is-art sharpening and the F-is-not-art sharpening*

As I hope is clear, the analysis of the link between the Mutt affair and the origins of the change in the practice that, drawing on Ásta's account, I provided in Section 4.2 cannot be the whole story, given the normative indeterminacy hypothesis. If it were the whole story, then, contrary to the hypothesis,  $F \sim \text{Art}$  (one of the propositions in Duchamp's paradox set) would be false, since Duchamp's submission of *Fountain* to the annual exhibition would have turned it into (determinate) art by conferring it the relevant property. The analysis can then only be part of the story. What is the rest of the story?

Consider the hypothesis that there are two sharpenings of 'is art':

S1: (For any  $x$ )  $x$  is art iff a subject (or group, or entity) in authority or with standing,  $S$ , conferred the property of being art onto  $x$  on the basis of  $S$ ' attempt to track some base property.

S2: (For any  $x$ )  $x$  is art iff  $x$  possesses property  $P$ .

As I understand these sharpenings, S1 corresponds to the sharpening on which *Fountain* is art (the *F-is-art* sharpening), while S2 corresponds to the sharpening on which *Fountain* is not art (the *F-is-not-art* sharpening). I want to remain as neutral as possible on what  $P$ , in S2, stands for, since there are several options and I need not commit to one of them to flesh out the normative indeterminacy hypothesis. However, to make my proposal as clear as possible, it is instructive to use (without endorsing) one of these options as an example of what  $P$  might stand for, namely, what I take to be one of the simplest options available: that  $P$  stands for the property of *belonging to at least one of the arts* (equivalently: for the property of being a work in at least one of the arts).<sup>144</sup> The reason that it is plausible that *Fountain* does not possess this property is that it clearly does not seem to belong to the art of sculpture or any other art, old or new.<sup>145</sup> Let us then assume, for the sake of argument, that  $P$  in S2 stands for the property of belonging to at least one of the arts, and that *Fountain* does not possess this property.

To return to the Mutt affair in light of sharpening S1, note that, given the analysis provided in Section 4.2, any member of the American Society of Independent Artists was a subject in authority. Moreover, note that, also given that analysis, those who actually submitted their work to the annual exhibition thereby conferred the property of being art onto it. Thus, on S1, their work is art. But which base property were they attempting to track in their conferrals? Consider those members of the Society who submitted paintings to the annual

---

<sup>144</sup> Another simple option here would be one that appeals to aesthetic properties. See, for example, Zangwill's [2007] aesthetic theory of art (a more sophisticated version than Beardsley's) for an example of what  $P$ , on this option, might stand for.

<sup>145</sup> *Pace* Lopes, who (as already discussed) thinks that, despite appearances to the contrary, it does belong to an art.

exhibition (the majority of the submissions). On the hypothesis that I am proposing, the property that, consciously or unconsciously, they were attempting to track in their conferrals was the property of being a painting. Similarly, the property that those members who submitted sculptures were, consciously or unconsciously, attempting to track in their conferrals was the property of being a sculpture. Thus, the former and latter members were attempting to track the base property of belonging to at least one of the arts—that is, *they were attempting to track P in S2*.

Suppose now, as is eminently plausible, that they *succeeded* in their attempts at tracking that base property. Then their works are art on all sharpenings, and so (determinately) art. Why? Because—letting S stand for any of those members and x for S' submitted work—if the base property that S was attempting to track in the conferral was the property of belonging to at least one of the arts, and S' attempt at tracking this property succeeded, then x belongs to at least one of the arts, in which case the proposition *x is art* is true on S1 and S2.

We are now in a position to understand what Duchamp (disguised as Mutt) did, at least according to the hypothesis that I am proposing here. In conferring the property of being art onto *Fountain* by submitting it to the annual exhibition, he was (consciously) attempting to track the base property of *not belonging to any of the arts*. In doing so, he showed that S1 and S2 could come apart if a subject in authority (or with standing) attempted to track—and succeeded at tracking—a property that, according to S2, is necessary and sufficient for *not* being art. Since *Fountain* is art on S1 but not on S2, the upshot is that it is indeterminate whether it is art—or, put another way, *Fountain* is indeterminate art.

On this way of fleshing out the normative indeterminacy hypothesis, then, one might say that S1 and S2 were *entangled* before Duchamp. This means that, before Duchamp, all those who had the authority or standing to confer the property of being art were, consciously or unconsciously, attempting to track—and succeeded at tracking (at least largely)—P in S2. Duchamp's insight was that S1 and S2 could be disentangled.



## 5. A SOLUTION TO THE DUCHAMP'S BOX PROBLEM

In Sections 3–4, I introduced the bare-bones of the normative indeterminacy hypothesis—as an unhappy-face solution to Duchamp's paradox—and fleshed out the hypothesis. But how can the normative indeterminacy hypothesis be established? Here is one obvious way to do it, given the preceding discussion: show that (i) the normative indeterminacy hypothesis yields a solution to the Duchamp's box problem, and that (ii) this hypothesis provides the best solution to the problem.

Now, because the Duchamp's box problem is a brand-new problem (despite being made up of some old problems), there is no independent literature on it. So, I am not in a position to convincingly establish (ii) in this dissertation, since that would involve comparing my solution to other solutions to the problem and arguing that mine is a better solution than each of those. However, I *am* in a position to convincingly establish (i) in this dissertation. This would be a significant result for obvious reasons: because the Duchamp's box problem is the problem for which we have been seeking a solution in this chapter, and because it is a hard problem. But there are other, perhaps less obvious reasons why the result would be significant: by establishing (i), I would thereby be forcing proponents of happy-face solutions to Duchamp's paradox *to provide a solution to the Duchamp's box problem*. Why? Because if the normative indeterminacy hypothesis entails that there is no happy-face solution to Duchamp's paradox (as I argued in Section 3) and if the hypothesis yields a solution to the Duchamp's box problem (as I shall argue in this section), then proponents of happy-face solutions to the paradox would have to provide a solution, and indeed a solution that is at least as good as mine, to the Duchamp's box problem—*on pain of having an explanatorily weaker solution to the paradox than mine (other things being equal)*.

In what follows, then, I argue that the normative indeterminacy hypothesis yields a solution to the Duchamp's box problem. In other words, I argue that the hypothesis yields a unified solution to this set of seven problems (the latter six of which are organised into pairs):

*Duchamp's paradox*

*The theoretical problem:* Is *Fountain* art? 

---

*The theoretical meta-problem:* Why does *Fountain* pose theoretical puzzlement, namely, puzzlement about whether it is art? *First Pair* 

---

*The problem of the change in the practice:* When and how did the radical post-*Fountain* change in the practice of treating things as art exactly occur? 

---

*The problem of the two avant-gardes via Dada:* What did the neo-avant-garde do—through Conceptualism—that had not already been done, to a large extent, by the historical avant-garde through Dada? *Second Pair* 

---

*The practical problem:* How, if at all, should *Fountain* be appreciated as art? 

---

*The practical meta-problem:* Why does *Fountain* pose practical puzzlement, namely, puzzlement about how, if at all, it should be appreciated as art? *Third Pair*

I have already argued that the normative indeterminacy hypothesis yields a solution to Duchamp's paradox. The question now is whether it yields a *unified* solution to it and the other problems. My argument for the conclusion that it does proceeds in several steps. In Section 5.1, I argue that the hypothesis yields a unified solution to the theoretical problem and the theoretical meta-problem—that is, to the First Pair of Problems. In Section 5.2, I return to a question left hanging in Chapter 2, the Action Under Normative Indeterminacy Question, and expand on what I said there about the features of action under normative indeterminacy. In Sections 5.3–4, partially drawing on this account of action under normative indeterminacy, I argue that the hypothesis yields a unified solution to the Second Pair of Problems (i.e., the problem of the change in the practice and the problem of the two avant-gardes via Dada) and a unified solution to the Third Pair of Problems (i.e., the practical problem and the practical meta-problem).

### 5.1. *A unified solution to the First Pair of Problems*

Given the solution to Duchamp's paradox that the normative indeterminacy hypothesis yields, a solution to the theoretical problem straightforwardly follows. Here are, once again, the propositions in the paradox set:

*Do Not TAA*: (For any  $x$ ) One must not: treat  $x$  as art if  $x$  is not art.

*F-May-be-TAA*: *Fountain* may be treated as art.

*F $\sim$ Art*: *Fountain* is not art.

On the unhappy-face solution to the paradox that the hypothesis yields, *F $\sim$ Art* is indeterminate. Does the hypothesis also yield a solution to the theoretical meta-problem, the problem of explaining why *Fountain* poses theoretical puzzlement (puzzlement about whether it is art)? To see that it does, recall that in Section 4, when fleshing out the normative indeterminacy hypothesis, I hypothesised that there are two sharpenings of 'is art':

S1: (For any  $x$ )  $x$  is art iff a subject (or group, or entity) in authority or with standing,  $S$ , conferred the property of being art onto  $x$  on the basis of  $S$ ' attempt to track some base property.

S2: (For any  $x$ )  $x$  is art iff  $x$  possesses property  $P$ .

Based on this, I argued that, in submitting *Fountain* to the Society's annual exhibition, Duchamp discovered that S1 and S2 could come apart if, in the conferral of the property of being art, a subject in authority or with standing attempts to track—and succeeds at tracking—a property that, given S2 alone, is necessary and sufficient for *not* being art. *This* is why *Fountain* poses theoretical puzzlement, and what, at the same time, makes it indeterminate whether *Fountain* is art—a *unified* solution to the theoretical meta-problem and the theoretical problem.

## 5.2. *Action under (unsuspected) normative indeterminacy*

Given the normative indeterminacy hypothesis, *Fountain* gives rise to normative indeterminacy, namely, indeterminacy about whether it ought to be treated as art and indeterminacy about whether it ought not to be treated as art. Suppose that one is in a choice situation in which it is indeterminate whether one ought to treat *Fountain* as art and indeterminate whether one ought not to treat *Fountain* as art. Then, given the view defended in Chapters 1–2, one faces a superhard choice. What are the features of one's action when one is acting in such a choice situation? This is an instance of the Action Under Normative Indeterminacy Question, which was left hanging in Chapter 2:

*Action Under Normative Indeterminacy Question:*

What are the features of one's action when one is acting in a choice situation such that it is indeterminate whether one ought to  $\varphi$  and indeterminate whether one ought not to  $\varphi$ ? In other words, what are the features of action under normative indeterminacy?

In Chapter 2, we saw that this question comes in two versions. On one version of it, the question asks what the features of one's action are when one is acting, *and knows or at least believes that one is acting*, under normative indeterminacy. I call this action under *suspected* normative indeterminacy (since, at the very least, one believes that one is acting under normative indeterminacy). On the other version of it, the question asks what the features of one's action are when one is acting, *and does not know or even believe that one is acting*, under normative indeterminacy. I call this action under *unsuspected* normative indeterminacy (since one does not even believe that one is acting under normative indeterminacy).

In what follows, I spell out five features of action under unsuspected normative indeterminacy (§5.2.1), and then use them to make sense of the Mutt affair (§5.2.2), since on my view—as should be clear, given the normative indeterminacy hypothesis—the board of

directors of the American Society of Independent Artists faced a superhard choice when they received *Fountain*.

#### 5.2.1. *Five features of action under unsuspected normative indeterminacy*

To be more precise about the choice situation in which we are interested, suppose that one is a minimally rational agent who faces, and believes that one faces, a choice situation in which  $\varphi$ ing is permissible iff  $p$ , and that  $\varphi$ ing is required iff  $p$ . If it is indeterminate whether  $p$ , then it is indeterminate whether one ought to  $\varphi$  and indeterminate whether one ought not to  $\varphi$ . Suppose that it is indeed indeterminate whether  $p$ , and that one does not know or believe that it is indeterminate whether  $p$ . Then one is acting under normative indeterminacy without knowing or believing that one is acting under it, what I have called action under unsuspected normative indeterminacy. This is the choice situation in which we are interested.

As I argued in Chapter 2, one key feature of action under normative indeterminacy—whether suspected or unsuspected—is that the normative indeterminacy is *inescapable*: no matter what one does, one cannot escape it. More precisely, given that one's choice situation is such that it is indeterminate whether one ought to  $\varphi$  and indeterminate whether one ought not to  $\varphi$ , *one is forced*—to put it in supervaluationist terms—to *act in accordance with some*, but in *conflict* with other, (admissible) *sharpenings of the normative indeterminacy*. In other words, no matter which course of action one takes,  $\varphi$ ing or not  $\varphi$ ing, one cannot escape acting under some, but against other, sharpenings of the normative indeterminacy. I call this the *inescapability feature* of action under (suspected and unsuspected) normative indeterminacy.

Having identified this key feature, we can ask: what happens when one acts, without knowing or believing that one is acting, under some but against other sharpenings of the normative indeterminacy? I want to suggest that four things happen, at least typically. *First*, one makes a *judgment* about an indeterminate matter without knowing or believing that one is making a judgment about an indeterminate matter. To see why I think this, note that, by  $\varphi$ ing,

one is acting under the sharpenings on which one *ought* to  $\varphi$ , so that the best explanation for one's behaviour—given that one believes that  $\varphi$ ing is required iff  $p$ —is that one made the judgment that  $p$ . Similarly, by not  $\varphi$ ing, one is acting under the sharpenings on which one ought *not* to  $\varphi$ , so that the best explanation for one's behaviour—given that one believes that  $\varphi$ ing is permissible iff  $p$ —is that one made the judgment that not- $p$ . Thus, I take it, when one cannot escape acting under some, but against other, sharpenings of the normative indeterminacy, one either judges that  $p$  or judges that not- $p$ , where it is indeterminate whether  $p$  and one does not know or believe this. Let us call this the *judgment feature* of action under unsuspected normative indeterminacy.

*Second*, when one makes the judgment that  $p$  or the judgment that not- $p$ —and thereby forms the corresponding belief—one is (rationally) *committed* to acting in a certain way. To see this, suppose that, faced with our choice situation, one judges, and thereby forms the belief, that  $p$ . Then, since one also believes that  $\varphi$ ing is required iff  $p$ , one is committed to  $\varphi$ ing, and thereby committed to acting under the sharpenings on which one ought to  $\varphi$  and against the sharpenings on which one ought not to  $\varphi$ . Similarly, if one judges, and so believes, that not- $p$ , then (since one also believes that  $\varphi$ ing is permissible iff  $p$ ) one is committed to not  $\varphi$ ing, and thereby committed to acting under the sharpenings on which one ought not to  $\varphi$  (and against the sharpenings on which one ought to  $\varphi$ ). Call this the *commitment feature* of action under unsuspected normative indeterminacy.

Notice that being committed to acting in a certain way holds regardless of whether one *believes* that one is so committed—that is, even if one does not believe that one is committed to acting in that way, one is still committed to so act. To bring this out, consider two simple (non-indeterminacy-related) cases due to Sam Shpall [2014: 149-50], the first of which is an example of what one might call *theoretical* commitment and the second an example of what one might call *practical* commitment:

*Bible*: Adam believes that everything that the Bible says is true, and he also believes that the Bible says that the world was created in six days. Then he is (rationally) committed to believing that the world was created in six days, even if he does not believe that the world was created in six days.

*Green jobs*: In the State of the Union address, the US president pledged to spend ten billion dollars on green jobs. Then the US president is (morally) committed to spending ten billion dollars on green jobs, even if the president does not want or intend to spend ten billion dollars on green jobs.

These two cases suggest, as Shpall [2014: 153] points out, that for one to be committed to something—be it to believing a certain proposition, or to acting in a certain way—is for one to stand in a *normative*, rather than psychological, relation to that thing. Why? Because one stands in that relation regardless of whether one believes to stand in it as well as regardless of whether one wants or intends to act on one's commitment, and the relation puts pressure on one to believe, or act on, that to which one is committed.

*Third*, and consistently with commitment being a normative relation, I take it that, when one makes the judgment/forms the belief that  $p$ , or that not- $p$ , and is thereby committed to acting in a certain way, one *recognises* (given the assumption that one is a minimally rational agent) that one is committed to acting in that way and acts accordingly. To put it another way, even though one need not believe that one is committed to acting in a certain way for one to be in fact committed to acting in that way, I take it that at least typically, when acting under unsuspected normative indeterminacy, one believes that one is so committed and acts in that way. Call this the *commitment recognition feature* of action under unsuspected normative indeterminacy.

*Fourth* (and finally), when one recognises that one is committed to acting in a certain way—say, to  $\varphi$ ing—and acts in that way, thereby acting under the sharpenings on which one ought to  $\varphi$ , this does not eliminate the indeterminacy about whether one ought/ought not to

φ. In other words, the normative indeterminacy does not go away—it *persists*. Let us call this the *persistence feature*. Why is this a feature of action under unsuspected normative indeterminacy? One obvious reason is that being committed to acting in a certain way is distinct from being required to acting in that way, so that acting on one’s commitment does not change what one is required to do. Consider again the Bible and Green Jobs cases.

Adam is (rationally) *committed* to believing that the world was created in six days, but is he (rationally) *required* to believing it? At least on a dominant view of what one might call *theoretical* requirements (the so-called ‘wide-scope’ view of rational requirements<sup>146</sup>), he is not. Rather, he is required *either* to believing that the world was created in six days *or* to giving up his belief that everything that the Bible says is true—and this seems right. In that case, though, there is a clear distinction between what Adam is committed to believing and what he is required to believe [Shpall 2014: 149]. Accordingly, there is a distinction (to use my terminology) between theoretical commitments and theoretical requirements. If this is right, there is presumably also a distinction between practical commitments and practical requirements. Here is a case for it. In the Green Jobs case, the US president pledged to spend, and is thereby (morally) committed to spending, ten billion dollars on green jobs. Suppose, however, that soon after the president’s pledge a catastrophic natural disaster (morally) requires that the ten billion dollars be diverted to help those affected by it [Shpall 2014: 150]. Then the president would be required to violate that commitment, in which case practical commitments and practical requirements are distinct from each other.

### 5.2.2. *The Mutt affair, revisited—yet again*

Let us roll back the tape, as it were, to the first few days of April 1917. The board of directors have just received *Fountain*, and a question arises: ought they to treat *Fountain* as art (viz., by

---

<sup>146</sup> See especially Broome [2007]



exhibiting it at the annual show), or ought they not to do it (viz., by not exhibiting it)? Let us assume that they believe that treating *Fountain* as art is permissible iff *Fountain* is art, and that they believe that treating *Fountain* as art is required iff *Fountain* is art. In other words, assume that they believe the TAA Norms to be in force for them. (As I suggested in Section 2, this assumption is not only plausible but needed to make sense of the Mutt affair.<sup>147</sup>) Furthermore, let us assume, plausibly, that they believe that there is a knowable fact of the matter about whether *Fountain* is art. Given these plausible assumptions and the normative indeterminacy hypothesis, they are acting under normative indeterminacy without knowing or believing that they are acting under it. Using the features of action under unsuspected normative indeterminacy to understand the Mutt affair, the following picture emerges.

#### 5.2.2.1. *The inescapability feature*

On the normative indeterminacy hypothesis, there are two sharpenings of ‘is art’:

S1: (For any  $x$ )  $x$  is art iff a subject (or group, or entity) in authority or with standing,  $S$ , conferred the property of being art onto  $x$  on the basis of  $S$ ’ attempt to track some base property.

S2: (For any  $x$ )  $x$  is art iff  $x$  possesses property  $P$ .

Given S1 together with the TAA Norms, the board of directors *ought* to treat *Fountain* as art. Given S2 together with the TAA Norms, the board of directors ought *not* to treat *Fountain* as art. No matter what they do then, whether they treat *Fountain* as art or whether they do not treat it as art, they cannot escape acting under one sharpening, but against the other sharpening, of the normative indeterminacy.

---

<sup>147</sup> To be more precise here: I argued in Section 2 that *Do Not TAA* is necessary to make sense of the behaviour of the board of directors. Now that Do TAA was also introduced I can add this: postulating Do TAA allows to make sense of why they exhibited the other submitted works.

#### 5.2.2.2. *The judgment feature*

After some disagreement about what to do, the board of directors decide that a majority vote be taken (among the directors present on the occasion) about whether to exhibit *Fountain* or not. In the process, each of the directors makes a judgment about whether *Fountain* is art, without knowing or believing that they are making a judgment about an indeterminate matter, and each of them forms the corresponding belief. The majority of the directors judge, and thereby believe, that *Fountain* is not art, and so the board's belief aligns with the belief of the majority.

#### 5.2.2.3. *The commitment & commitment recognition features*

Since the board believe that treating *Fountain* as art is permissible iff *Fountain* is art, and also believe that *Fountain* is not art, they are (rationally) committed to acting in a certain way, namely, to not treating *Fountain* as art. Recognising their commitment, they act on it by not treating *Fountain* as art, thereby acting under the sharpening of the normative indeterminacy on which they ought not to treat *Fountain* as art and against the sharpening on which they ought to treat *Fountain* as art. Later, they issue this revealing statement (already quoted in Section 2.1) to the press: "The *Fountain* may be a very useful object in its place, but its place is not an art exhibition and it is by no definition a work of art."<sup>148</sup>

#### 5.2.2.4. *The persistence feature*

Finally, by recognising their commitment and thereby acting under the sharpening on which they ought not to treat *Fountain* as art, and against the sharpening on which they ought to treat *Fountain* as art, the board did not eliminate the indeterminacy about whether *Fountain*

---

<sup>148</sup> See n. 101 above.

ought/ought not to be treated as art. That is, their acting under that sharpening did not make the normative indeterminacy to which *Fountain* gives rise go away—the normative indeterminacy persists.

### 5.3. *A unified solution to the Second Pair of Problems*

Recall the Second Pair of Problems:

*The problem of the change in the practice:* When and how did the radical post-*Fountain* change in the practice of treating things as art exactly occur?

*The problem of the two avant-gardes via Dada:* What did the neo-avant-garde do—through Conceptualism—that had not already been done, to a large extent, by the historical avant-garde through Dada?

In order to spell out a unified solution to these problems, consider first the second half of the problem of the change in the practice, the question of *how* the change in the practice of treating things as art exactly occurred. Given the normative indeterminacy hypothesis and the features of action under unsuspected normative indeterminacy identified in Section 5.2, Duchamp's discovery that sharpenings S1 and S2 could come apart and the board's action of not treating *Fountain* as art constituted a *set of instructions on how to change the practice*:

Instruction 1: In your act of conferral onto  $x$ , show that, as a subject (or group, or entity) in authority or with standing, you can make S1 and S2 come apart by successfully tracking a base property that, on S2, is necessary and sufficient for not being art; and

Instruction 2: Have a sufficiently large group of participants in the practice judge, and form the belief, that  $x$  is art, and thereby be committed (given a belief that the TAA Norms are in force) to treating  $x$  as art, recognise their commitment and act on it. That is, have the group do the opposite of what the board of directors did: act under the

sharpening on which  $x$  ought to be treated as art and against the sharpening on which  $x$  ought not to be treated as art.

Importantly,  $x$  could stand for *any* work that is art on S1 but not on S2, *Fountain* being merely one example—though the paradigmatic one—of such a work. Because the change in the practice of treating things as art did occur (it is an empirical fact), what the normative indeterminacy hypothesis predicts is that *at some point in time* members of a sufficiently large group triggered the change by carrying out those instructions. The question now is *when*—that is, when did the change in the practice exactly occur?

A plausible empirical hypothesis is that this occurred when, in the process of reception of Duchamp's readymades and especially *Fountain*, members of the neo-avant-garde—namely, Conceptualists—triggered the change by carrying out those instructions. In carrying out the instructions, they did something that the historical avant-garde through Dada had not done. As a matter of historical fact, what the historical avant-garde did through Dada was to supply the instructions on how to change the practice, not to carry them out. To a significant extent, this reconciles some of the central theses of the Bürgerian and Fosterian solutions to the problem of the two avant-gardes via Dada. One central Bürgerian thesis is that the project of the neo-avant-garde through Conceptualism was derivative of the project of Dada, and one central Fosterian thesis is that the neo-avant-garde enacted the project of Dada for the first time. On the solution that I propose here, there is something to both theses, since what Conceptualists did was to carry out the instructions supplied by Dada (hence the derivativeness) and in doing so they changed the practice (hence the enactment).

This constitutes a *unified* solution to the Second Pair of Problems because the solution to the problem of the two-avant-gardes via Dada is also a solution to the problem of *when* (viz., in the neo-avant-garde period) and *how* (viz., by carrying out the instructions) the change in the practice occurred.

#### 5.4. *A unified solution to the Third Pair of Problems*

One implication of the discussion in Sections 5.2–3 is that while members of the neo-avant-garde changed the practice through their many actions under unsuspected normative indeterminacy—more specifically, through their actions under the sharpening on which  $x$  ought to be treated as art, where  $x$  stands for any work that is art on S1 but not on S2—the normative indeterminacy to which  $x$  gives rise *persists*. The paradigmatic example of  $x$ , of course, is *Fountain*. As such, even though today *Fountain* is exhibited at the Philadelphia Museum of Art (for example<sup>149</sup>) and so treated as art, we should expect—if my hypothesis is true—the normative indeterminacy to which *Fountain* gives rise to manifest itself in some way. Throughout this chapter, I have time and again discussed *one* way in which it manifested itself, namely, when in 1917 the board of directors received *Fountain* and had to deliberate whether to exhibit it. Are there other ways in which the normative indeterminacy might manifest itself?

To see that there are, recall that (as mentioned in Section 2.5) this simple claim is true: if one appreciates  $x$  as art, then one is treating  $x$  as art. This entails (as we saw in Section 2.5) that if one appreciates  $x$  as art, but  $x$  is not art, then Do Not TAA is violated (and so are, by implication, the TAA Norms), since, according to this norm, one ought not to treat as art that which is not art. In fact, Glasses—introduced in Section 3 as evidence for Do Not TAA—should be understood as a case in which my glasses are appreciated as (but are not) art, and so are treated as art:

*Glasses*: Suppose that, while visiting one of the rooms dedicated to avant-garde works of the 1960s at Tate Modern’s permanent collection, I accidentally drop my glasses, without myself or anyone else around me noticing it, and proceed to the next room without them. As other museum-goers enter the room I was in, they start treating my glasses as art—say, by taking

---

<sup>149</sup> As mentioned in n. 89, this is just one of the versions of *Fountain* out there.

photographs of them as they lie on the floor, ascribing meaning to them (in order to retrieve the artist's intention), and so on.

By Do Not TAA, that simple claim implies (together an uncontroversial claim) that there is a *biconditional* norm governing appreciation within the practice of treating things as art. Why? Because, as we have just seen, if one appreciates  $x$  as art, but  $x$  is not art, then, by the Do Not TAA, one is doing something that one ought not to do (viz., treating as art that which is not art). So, if  $x$  is not art, one ought not to appreciate it as art (equivalently: if one may appreciate  $x$  as art, then  $x$  is art). Since it is uncontroversial that if  $x$  is art, then one may appreciate  $x$  as art, it follows that this norm governs appreciation within the practice:

*May AAA:* (For any  $x$ ) One may: appreciate  $x$  as art iff  $x$  is art.

Given the normative indeterminacy hypothesis and this norm, it follows that it is indeterminate whether *Fountain* may be appreciated as art—another way in which the normative indeterminacy to which *Fountain* gives rise manifests itself.

With that in mind, recall the Third Pair of Problems:

*The practical problem:* How, if at all, should *Fountain* be appreciated as art?

*The practical meta-problem:* Why does *Fountain* pose practical puzzlement, namely, puzzlement about how, if at all, it should be appreciated as art?

As we saw in Section 2, the practical problem breaks into two questions:

Qa: May *Fountain* be appreciated as art? In other words, is it appropriate to appreciate *Fountain* as art?

Qb: If so, how should *Fountain* be appreciated as art? In other words, what is the appropriate way to appreciate *Fountain* as art?

For those who propose to solve the problem by answering ‘Yes’ to Qa, Qb arises. For those who propose to solve it by answering ‘No’ to Qa, Qb does not arise.

Here then is how the normative indeterminacy hypothesis provides a unified solution to the Third Pair of Problems. On S2 (the sharpening on which *Fountain* is not art), *Fountain* should *not* be appreciated as art—hence why it poses puzzlement about *whether* it may be appreciated as art at all. On S1 (the sharpening on which *Fountain* is art), *Fountain* may be appreciated as art, but Qb arises because *Fountain* does not possess property P in S2—hence why it poses puzzlement about *how* it should be appreciated as art. This is why *Fountain* poses practical puzzlement, puzzlement about how, if at all, it should be appreciated as art. At the same time, this solution to the practical meta-problem is also a solution to the practical problem, since it implies that it is indeterminate whether it is permissible (appropriate) to appreciate *Fountain* as art, and so indeterminate whether there is a required (appropriate) way to appreciate it as art.

### 5.5. *A unified solution to all problems*

Let us take stock. In Sections 3–4, I put forward the normative indeterminacy hypothesis as a solution to Duchamp’s paradox. In Section 5.1, I argued, by appealing to sharpenings S1 and S2, that the hypothesis yields a unified solution to the First Pair of Problems. In Section 5.3, I argued, by appealing to the features of action under unsuspected normative indeterminacy (as spelt out in Section 5.2), that it yields a unified solution to the Second Pair of Problems. Finally, in Section 5.4, I argued, by appealing to the persistence of normative indeterminacy, that it yields a unified solution to the Third Pair of Problems. Why do these solutions—to Duchamp’s paradox and the three pairs of problems—constitute a *single, unified* solution to all problems? Because there is a common phenomenon to which they appeal: the indeterminacy of ‘is art’.

## 6. AN ALTERNATIVE INTERPRETATION OF THE HYPOTHESIS

What did Duchamp discover in the 1910s through his paradox? According to the normative indeterminacy hypothesis, he discovered that two sharpenings of ‘is art’ could come apart, since *Fountain* is art on one of them (S1) but not the other (S2). This is one of the core claims of the normative indeterminacy hypothesis (the other being that the TAA Norms govern the practice of treating things as art). In the previous sections of this chapter, I interpreted this core claim in a particular way: that the relevant sharpenings (S1 and S2) are still sharpenings of ‘is art’. Given this interpretation, it is indeterminate whether *Fountain*—and indeed any work that is art on S1 but not on S2—is art. There is, however, an alternative interpretation of that core claim: that one of the relevant sharpenings is no longer a sharpening of ‘is art’, namely, S2. Given this alternative interpretation, it is no longer indeterminate whether *Fountain* is art: *Fountain* is now (determinately) art. How to make sense of this interpretation of the normative indeterminacy hypothesis?

According to the alternative interpretation of the hypothesis, *Fountain* gave rise to normative indeterminacy at least in the 1910s (since S1 and S2 were sharpenings of ‘is art’ then and the TAA Norms govern the practice), but now it is (determinately) permissible to treat it as art. Suppose that the source of the indeterminacy of ‘is art’ is purely semantic. Given this, one way of making sense of the alternative interpretation is as follows. As used before the 1910s, the predicate ‘is art’ means either S1 or S2. But the linguistic practices of those using ‘is art’ before the 1910s leave it indeterminate whether they meant S1 or S2 by it, since they never had to apply the predicate in situations in which the two sharpenings came apart—hence the theoretical and practical puzzlement posed by *Fountain* in the 1910s, and in fact hence Duchamp’s paradox.<sup>150</sup> Thus, in carrying out the instructions supplied by Dada, members of

---

<sup>150</sup> Compare Hartry Field [1973] on the use of ‘mass’ before Einstein’s discovery of the theory of special relativity.



the neo-avant-garde changed not only the practice of treating things as art but also the meaning of ‘is art’, namely, by eliminating S2. They eliminated S2 because—given their systematic actions *in accordance* with the sharpening on which it is permissible, and *in conflict* with the one on which it is not permissible, to treat works like *Fountain* as art—our linguistic practices from the mid-1970s (the end of the neo-avant-garde) onwards no longer leave it indeterminate what we mean by ‘is art’: we mean S1 and S1 only.

There is something undeniably attractive about this interpretation of the hypothesis. Nevertheless, endorsing it appears to come at a price: the interpretation appears unable to yield a solution to the Duchamp’s box problem. Why? Because while it can explain why *Fountain* posed theoretical and practical puzzlement in the 1910s and the decades before members of the neo-avant-garde changed the meaning of ‘is art’, it appears unable to explain why *Fountain* poses theoretical and practical puzzlement (puzzlement about whether it is art and about how, if at all, it should be appreciated as art) *to this day*. In other words, it appears unable to yield a solution to the meta-problems, since it apparently predicts that the theoretical and practical puzzlement posed by *Fountain* should have gone away by now. Given that we need a hypothesis that yields a solution to the meta-problems, this is a price that we should not be willing to pay—hence why I interpret the normative indeterminacy hypothesis in the way I do.

## 7. CONCLUSION

In this chapter, drawing and expanding on the connection between choice and indeterminacy established in Chapter 2, I developed a framework with which to solve the Duchamp’s box problem, the neglected problem of finding a unified solution to a set of hard problems that have *Fountain* at their core. Whether this framework—the normative indeterminacy hypothesis—yields the best solution to the Duchamp’s box problem is an open question. What I take myself to have shown is that it yields a solution to it, and this is no small feat. The normative indeterminacy hypothesis is worth serious consideration.

## Conclusion

In this dissertation, I developed solutions to three problems in value theory, broadly understood: the problem of superhard comparisons, (what might be called) the choice under moral vagueness problem, and the Duchamp's box problem. To that end, I adopted the same strategy throughout: to develop accounts of decision-making under vagueness (or indeterminacy, more generally), conditional on some leading theories of vagueness. More specifically, to develop solutions to those problems, each chapter of the dissertation developed an answer—or at least a partial answer—to at least one of the following questions:

*Choice Under Vague Betterness Question:*

In a choice situation between a pair of options,  $x$  and  $y$ , what is it that one rationally ought to do if one knows that it is vague whether  $x$  is better than  $y$ ?

*Choice Under Moral Vagueness Question:*

What is it that one rationally ought to do if one knows that it is (determinately) prudentially better to  $\varphi$ , but that it is vague whether it is morally permissible to  $\varphi$ ?

*Action Under Normative Indeterminacy Question:*

What are the features of one's action when one is acting in a choice situation such that it is indeterminate whether one ought to  $\varphi$  and indeterminate whether one ought not to  $\varphi$ ? In other words, what are the features of action under normative indeterminacy?

Chapter 1 developed a partial answer to the Choice Under Vague Betterness Question—namely, an account of decision-making under asymmetric vague betterness—in the context of the Make No Mistake argument against Trichotomy-denying views and for the new vagueness view of superhard comparisons and choices (conditional on any Must Not Believe It—entailing theory of vagueness).

Chapter 2 developed at least a partial answer to each of the three questions in the context of an argument with a twofold aim. One aim of this argument was to complete my defence of the new vagueness view, conditional on vagueness-as-absence-of-truth. The other aim of it was to show that while vagueness-as-absence-of-truth (together with TNB) and vagueness-as-ignorance (together with KNB) yield the same account of decision-making under asymmetric vague betterness—given the Make No Mistake argument—they yield different accounts of decision-making under symmetric vague betterness and this has implications for some debates in practical ethics.

Chapter 3 developed a partial answer to the Action Under Normative Indeterminacy Question (viz., an account of action under unsuspected normative indeterminacy) in the context of a solution to the Duchamp's box problem—the normative indeterminacy hypothesis, according to which 'is art' has two sharpenings (S1 and S2) and there are norms governing the practice of treating things as art (the TAA Norms).



## Bibliography

- Abell, Catharine [2012], 'Art: What It is and Why It Matters', *Philosophy and Phenomenological Research*, 85: 671-91.
- Alberro, Alexander [1999], 'Reconsidering Conceptual Art, 1966-1977', in Alexander Alberro and Blake Stimson (eds.), *Conceptual Art: A Critical Anthology*, Cambridge, MA: MIT Press, xvi-xxxvii.
- Allen, Colin, and Trestman, Michael [2016], 'Animal Consciousness', in Edward Zalta (ed.), *Stanford Encyclopedia of Philosophy*: <[plato.stanford.edu/entries/consciousness-animal/](http://plato.stanford.edu/entries/consciousness-animal/)>.
- Antony, Michael [2006a], 'Papineau on the Vagueness of Phenomenal Concepts', *Dialectica*, 60: 475-83.
- [2006b], 'Vagueness and the Metaphysics of Consciousness', *Philosophical Studies*, 128: 515-38.
- Ásta [2018], *Categories We Live By*, New York: Oxford University Press.
- Bales, Adam [2018], 'Indeterminate Permissibility and Choiceworthy Options', *Philosophical Studies*, 175: 1693-702.
- Barnes, Elizabeth [2010], 'Ontic Vagueness: A Guide for the Perplexed', *Noûs*, 44: 601-27.
- and Williams, Robert [2011], 'A Theory of Metaphysical Indeterminacy', *Oxford Studies in Metaphysics*, 6: 103-48.
- Barnett, David [2009], 'Is Vagueness Sui Generis?', *Australasian Journal of Philosophy*, 87: 5-34.
- [2011], 'Does Vagueness Exclude Knowledge?', *Philosophy and Phenomenological Research*, 82: 22-45.
- [MS] 'Vagueness and Rationality'.
- Beardsley, Monroe [1981], *Aesthetics*, 2<sup>nd</sup> ed., Indianapolis: Hackett.
- [1983] (2006), 'An Aesthetic Definition of Art', in Peter Lamarque and Stein Haugom Olsen (eds.), *Aesthetics and the Philosophy of Art*, Malden: Blackwell, 55-62.
- Binkley, Timothy [1970], 'Piece: Contra Aesthetics', *Journal of Aesthetics and Art Criticism*, 35: 265-77.
- Block, Ned [1992], 'Begging the Question Against Phenomenal Consciousness', *Behavioral and Brain Sciences*, 15: 205-6.
- [2002], 'The Harder Problem of Consciousness', *Journal of Philosophy*, 99: 391-425.
- Boonin, David [2003], *A Defense of Abortion*, Cambridge: Cambridge University Press.
- Bourdieu, Pierre [2017], *Manet: A Symbolic Revolution*, trans. Peter Collier and Margaret Rigaud-Drayton, Cambridge: Polity.
- Braithwaite, Victoria [2010], *Do Fish Feel Pain?*, Oxford: Oxford University Press.
- Broome, John [1997], 'Is Incommensurability Vagueness?', in Ruth Chang (ed.), *Incommensurability, Incomparability, and Practical Reasoning*, Cambridge, MA: Harvard University Press: 67-89.
- [1999], *Ethics out of Economics*, Cambridge: Cambridge University Press.
- [2004], *Weighing Lives*, Oxford: Oxford University Press.
- [2007], 'Wide and Narrow Scope', *Mind*, 116: 359-70.
- [2009], 'Reply to Rabinowicz', *Philosophical Issues*, 19: 412-7.

- Buchloh, Benjamin [1984], 'Theorizing the Avant-Garde', *Art in America*, 72.10: 19-20.
- [1986], 'The Primary Colors for the Second Time: A Paradigm Repetition of the Neo-Avant-Garde', *October*, 37: 41-52.
- [1990], 'Conceptual Art 1962-1969: From the Aesthetic of Administration to the Critique of Institutions', *October*, 55: 105-143.
- [1996], 'Three Conversations in 1985: Claes Oldenburg, Andy Warhol, Robert Morris', in Martha Buskirk and Mignon Nixon (eds.), *The Duchamp Effect*, Cambridge, MA: MIT Press, 33-54.
- [2000], *Neo-Avant-Garde and Culture Industry*, Cambridge, MA: MIT Press.
- Buchloh, Benjamin et al. [1996], 'Conceptual Art and the Reception of Duchamp', in Martha Buskirk and Mignon Nixon (eds.), *The Duchamp Effect*, Cambridge, MA: MIT Press, 205-224.
- Bürger, Peter [1984], *Theory of the Avant-Garde*, trans. Michael Shaw, Minneapolis: University of Minnesota Press.
- [2010], 'Avant-Garde and Neo-Avant-Garde: An Attempt to Answer Certain Critics of *Theory of the Avant-Garde*', *New Literary History*, 41: 695-715.
- Caie, Michael [2012], 'Belief and Indeterminacy', *Philosophical Review*, 121: 1-54.
- Camfield, William [1989], *Marcel Duchamp: Fountain*, Houston: Menil Foundation/Houston Fine Art Press.
- Carlson, Erik [2004], 'Broome's Argument Against Value Incomparability', *Utilitas*, 16: 220-4.
- [2013], 'Vagueness, Incomparability, and the Collapsing Principle', *Ethical Theory and Moral Practice*, 16: 449-63.
- Carroll, Noël [1986], 'Art and Interaction', *Journal of Aesthetics and Art Criticism*, 45: 57-68.
- [2004], 'Non-Perceptual Aesthetic Properties: Comments for James Shelley', *British Journal of Aesthetics*, 44: 413-23.
- Chalmers, David [1996], *The Conscious Mind*, New York: Oxford University Press.
- Chang, Ruth [2002a], 'The Possibility of Parity', *Ethics*, 112: 659-88.
- [2002b], *Making Comparisons Count*, London: Routledge.
- [2005], 'Parity, Interval Value and Choice', *Ethics*, 115: 331-50.
- [2012], 'Are Hard Choices Cases of Incomparability?', *Philosophical Issues* (supplement to *Noûs*) 22: 106-26.
- [2016], 'Comparativism: The Grounds of Rational Choice', in Errol Lord and Barry Maguire (eds.), *Weighing Reasons*, Oxford: Oxford University Press: 213-40.
- Cohen, Ted [1973], 'The Possibility of Art: Remarks on a Proposal by Dickie', *Philosophical Review*, 82: 69-82.
- Colb, Sherry F., and Dorf, Michael C. [2016], *Beating Hearts: Abortion and Animal Rights*, New York: Columbia University Press.
- Constantinescu, Cristian [2012], 'Value Incomparability and Indeterminacy', *Ethical Theory and Moral Practice*, 15: 57-70.
- [2014], 'Moral Vagueness: A Dilemma for Non-Naturalism', *Oxford Studies in Metaethics*, 9: 152-85.
- Costello, Diarmuid [2013], 'Kant and the Problem of Strong Non-Perceptual Art', *British Journal of Aesthetics*, 53: 277-98.
- Danto, Arthur [1964], 'The Artworld', *Journal of Philosophy*, 61: 571-84.

- [1981], *The Transfiguration of the Commonplace: A Philosophy of Art*, Cambridge, MA: Harvard University Press.
- Davies, David [2004], *Art As Performance*, Malden, MA: Blackwell.
- Davies, Stephen [1991], *Definitions of Art*, Ithaca: Cornell University Press.
- De Duve, Thierry [1996], *Kant After Duchamp*, Cambridge, MA: MIT Press.
- [2013a], ‘Pardon My French’, *Artforum*, 52.2: 246, 248-50, 252-3, 318.
- [2013b], ‘Don’t Shoot the Messenger’, *Artforum*, 52.3: 265-273, 312.
- [2014a], ‘Why Was Modernism Born in France?’, *Artforum*, 52.4: 190-7, 234-5.
- [2014b], ‘The Invention of Non-Art: A History’, *Artforum*, 52.6: 192-9, 238.
- [2014c], ‘The Invention of Non-Art: A Theory’, *Artforum*, 52.7: 270-5, 308-9.
- [2014d], ‘“This is art”: Anatomy of a Sentence’, *Artforum*, 52.8: 242, 245-6, 248-9, 282.
- De Sousa, Ronald [1974], ‘The Good and the True’, *Mind*, 83: 534-51.
- Dennett, Daniel [1995], ‘Animal Consciousness: What Matters and Why’, *Social Research*, 62: 691-710.
- Dickie, George [1969], ‘Defining Art’, *American Philosophical Quarterly*, 6: 253-6.
- [1971], *Aesthetics*, Indianapolis: Pegasus.
- [1974], *Art and the Aesthetic*, Ithaca: Cornell University Press.
- [1984], *The Art Circle*, New York: Haven Publications.
- [2000], ‘The Institutional Theory of Art’, in Noël Carroll (ed.), *Theories of Art Today*, Madison: University of Wisconsin Press, 93-108.
- Dodd, Julian [2016], ‘The Ontology of Conceptual Art: Against the Idea Idea’, in Julian Dodd (ed.), *Art, Mind, and Narrative: Themes from the Work of Peter Goldie*, Oxford: Oxford University Press.
- [2017], ‘What 4’33” is’, *Australasian Journal of Philosophy*, 96: 629-41.
- Dorr, Cian [2003], ‘Vagueness without Ignorance’, *Philosophical Perspectives*, 17: 83-113.
- Dorsey, Dale [2013], ‘Two Dualisms of Practical Reason’, *Oxford Studies in Metaethics*, 8: 114-39.
- Dougherty, Tom [2014], ‘Vague Value’, *Philosophy and Phenomenological Research*, 89: 352-72.
- Duchamp, Marcel [1912-20] (1975), ‘A L’Infini’, in Michel Sanouillet and Elmer Peterson (eds.), *The Essential Writings of Marcel Duchamp*, London: Thames & Hudson, 74-101.
- [1961] (1975), ‘Apropos of “Readymades”’, in Michel Sanouillet and Elmer Peterson (eds.), *The Essential Writings of Marcel Duchamp*, London: Thames & Hudson, 141-2.
- Dunaway, Billy [2017], ‘Ethical Vagueness and Practical Reasoning’, *Philosophical Quarterly*, 67: 38-60.
- Eklund, Matti [2011], ‘Being Metaphysically Unsettled’, *Oxford Studies in Metaphysics*, 6: 149-72.
- [2013], ‘Williams on the Normative Silence of Indeterminacy’, *Analysis*, 73: 265-71.
- Elga, Adam [2010], ‘Subjective Probabilities Should be Sharp’, *Philosophers’ Imprint*, 10: 1-11.
- Elson, Luke [2014a], ‘Heaps and Chains: Is the Chaining Argument for Parity a Sorites?’, *Ethics*, 124: 557-71.
- [2014b], ‘Borderline Cases and the Collapsing Principle’, *Utilitas*, 26: 51-60.
- Espinoza, Nicolas [2008], ‘The Small Improvement Argument’, *Synthese*, 165: 127-39.
- Field, Hartry [1973], ‘Theory Change and the Indeterminacy of Reference’, *Journal of Philosophy*, 70: 462-81.
- [2000], ‘Indeterminacy, Degree of Belief, and Excluded Middle’, *Noûs*, 34: 1-30.

- Filipovic, Elena [2016], *The Apparently Marginal Activities of Marcel Duchamp*, Cambridge, MA: MIT Press.
- Fine, Kit [1975], 'Vagueness, Truth and Logic', *Synthese*, 30: 265-300.
- Finer, Lawrence et al. (2005), 'Reasons U.S. Women have Abortions', *Perspectives on Sexual and Reproductive Health*, 37: 110-8.
- Foster, Hal [1996], 'What's Neo About the Neo-Avant-Garde?', in Martha Buskirk and Mignon Nixon (eds.), *The Duchamp Effect*, Cambridge, MA: MIT Press, 5-32.
- Fried, Michael [1996], *Manet's Modernism: Or, the Face of Painting in the 1860s*, Chicago: Chicago University Press.
- Friedman, Jane [2013a], 'Suspended Judgment', *Philosophical Studies*, 162: 165-81.
- [2013b], 'Rational Agnosticism and Degrees of Belief', *Oxford Studies in Epistemology*, 4: 57-81.
- [2017], 'Why Suspend Judging?', *Noûs*, 51: 302-26.
- Gaiger, Jason [2009], 'Dismantling the Frame: Site-Specific Art and Aesthetic Autonomy', *British Journal of Aesthetics*, 49: 43-58.
- Gaut, Berys [2000], "'Art" as a Cluster Concept', in Noël Carroll (ed.), *Theories of Art Today*, Madison: University of Wisconsin Press, 25-44.
- Gilbert, Margaret [2018], *Rights and Demands*, New York: Oxford University Press.
- Glackens, William et al. [1917], *Catalogue of the First Annual Exhibition of the Society of Independent Artists*, New York: Society of Independent Artists, Inc.
- Greenberg, Clement [1986], 'Avant-Garde Attitudes: New Art in the Sixties', in Clement Greenberg, *The Collected Essays and Criticism*, Vol. 4, Chicago: University of Chicago Press, 292-303.
- Greenough, Patrick [2003], 'Vagueness: A Minimal Theory', *Mind*, 112: 235-81.
- [2008], 'Indeterminate Truth', *Midwest Studies in Philosophy*, 32: 213-41.
- Goldie, Peter, and Schellekens, Elisabeth [2007], 'Introduction', in Peter Goldie and Elisabeth Schellekens (eds.), *Philosophy and Conceptual Art*, Oxford: Oxford University Press, ix-xxi.
- Griffin, James [1986], *Well-Being*, Oxford: Oxford University Press.
- Gustafsson, Johan [2013], 'Indeterminacy and the Small-Improvement Argument', *Utilitas*, 25: 433-45.
- Harman, Elizabeth [1999], 'Creation Ethics: The Moral Status of Early Fetuses and the Ethics of Abortion', *Philosophy & Public Affairs*, 28: 310-24.
- [2015], 'Morality Within the Realm of the Morally Permissible', *Oxford Studies in Normative Ethics*, 5: 221-44.
- [2016], 'Morally Permissible Moral Mistakes', *Ethics*, 126: 366-93.
- Hopkins, David [2004], *Dada and Surrealism*, Oxford: Oxford University Press.
- [2006], 'Re-Thinking the 'Duchamp Effect'', in Amelia Jones (ed.), *A Companion to Contemporary Art Since 1945*, Malden, MA: Blackwell, 145-63.
- [2016], 'New York Dada: From End to Beginning', in David Hopkins (ed.), *A Companion to Dada and Surrealism*, Malden, MA: Blackwell, 70-88.
- Hopkins, Robert [2007], 'Speaking Through Silence: Conceptual Art and Conversational Implicature', in Peter Goldie and Elisabeth Schellekens (eds.), *Philosophy and Conceptual Art*, Oxford: Oxford University Press, 51-67.
- Horgan, Terry et al. [2019], 'The Soritical Centipede', *Noûs*, 53: 491-510.



- Hursthouse, Rosalind [1991], 'Virtue Theory and Abortion', *Philosophy & Public Affairs*, 20: 223-46.
- [2000], *Ethics, Humans and Other Animals*, London: Routledge.
- Jago, Mark [2013], 'The Problem with Truthmaker-Gap Epistemicism', *Thought*, 320-29.
- Kaczor, Christopher [2015] (2010), *The Ethics of Abortion*, 2<sup>nd</sup> ed., New York: Routledge.
- Kearns, Stephen, and Magidor, Ofra [2008], 'Epistemicism about Vagueness and Metalinguistic Safety', *Philosophical Perspectives*, 22: 277-304.
- Keefe, Rosanna [2003], *Theories of Vagueness*, Cambridge: Cambridge University Press.
- Korsgaard, Christine [2018], *Fellow Creatures*, Oxford: Oxford University Press.
- Lee, Patrick [2010], *Abortion and Unborn Human Life*, 2<sup>nd</sup> ed., Washington, D.C.: The Catholic University of America Press.
- Levinson, Jerrold [1979], 'Defining Art Historically', *British Journal of Aesthetics*, 19: 232-50.
- [1989], 'Refining Art Historically', *Journal of Aesthetics and Art Criticism*, 47: 21-33.
- [1993], 'Extending Art Historically', *Journal of Aesthetics and Art Criticism*, 51: 411-23.
- Lewis, David [1987], *On the Plurality of Worlds*, Oxford: Blackwell.
- Lippard, Lucy, and Chandler, John [1971], 'The Dematerialization of Art', in Lucy Lippard, *Changing: Essays in Art Criticism*, New York: E. P. Dutton & Co.
- Little, Margaret [2008], 'Abortion and the Margins of Personhood', *Rutgers Law Journal*, 39: 331-48.
- Littlejohn, Clayton [2013], 'The Russellian Retreat', *Proceedings of the Aristotelian Society*, 113: 293-320.
- Lodder, Christina [1983], *Russian Constructivism*, New Haven: Yale University Press.
- Lopes, Dominic McIver [2007], 'Conceptual Art is Not What It Seems', in Peter Goldie and Elisabeth Schellekens (eds.), *Philosophy and Conceptual Art*, Oxford: Oxford University Press, 238-56.
- [2008], 'Nobody Needs a Theory of Art', *Journal of Philosophy*, 105: 109-27.
- [2014], *Beyond Art*, Oxford: Oxford University Press.
- Lord, Errol [2017], 'What You're Rationally Required to Do and What You Ought to Do (Are the Same Thing!)', *Mind*, 126: 1109-54.
- Mackenzie, Catriona [1992], 'Abortion and Embodiment', *Australasian Journal of Philosophy*, 70: 136-55.
- Mainardi, Patricia [1993], *The End of the Salon*, Cambridge: Cambridge University Press.
- Mandelbaum, Maurice [1965], 'Family Resemblances and Generalization Concerning the Arts', *American Philosophical Quarterly*, 2: 219-28.
- Manley, David [MS], 'Moral Realism and Semantic Plasticity'.
- Marquis, Don [1989], 'Why Abortion is Immoral', *Journal of Philosophy*, 86: 183-202.
- McElwee, Brian [2017], 'Supererogation Across Normative Domains', *Australasian Journal of Philosophy*, 95: 505-16.
- McGee, Vann, and McLaughlin, Brian [1995], 'Distinctions without a Difference', *Southern Journal of Philosophy*, 33: 203-51.
- McGinn, Colin [1982], *The Character of Mind*, New York: Oxford University Press.
- McMahan, Jeff [2002], *The Ethics of Killing*, New York: Oxford University Press.
- Naumann, Francis [1990], 'Marcel Duchamp: A Reconciliation of Opposites', in Rudolf Kuenzli and Francis Naumann (eds.), *Marcel Duchamp: Artist of the Century*, Cambridge, MA: MIT Press, 20-40.

- [1994], *New York Dada: 1915-23*, New York: Harry N. Abrams Publishers.
- [2012], *The Recurrent, Haunting Ghost: Essays on the Art, Life and Legacy of Marcel Duchamp*, New York: Readymade Press.
- Nesbit, Molly, and Sawelson-Gorse, Naomi [1996], ‘Concept of Nothing: New Notes by Marcel Duchamp and Walter Arensberg’, in Martha Buskirk and Mignon Nixon (eds.), *The Duchamp Effect*, Cambridge, MA: MIT Press, 131-75.
- Norton, Louise [1917], ‘Buddha of the Bathroom’, *The Blind Man*, 2: 5-6.
- O’Rourke, Joshua [MS], ‘Phenomenal Consciousness Must be Sharp’.
- Overall, Christine [1987], *Ethics and Human Reproduction: A Feminist Analysis*, London: Allen & Unwin.
- Overall, Christine [2015], ‘Rethinking Abortion, Ectogenesis, and Fetal Death’, *Journal of Social Philosophy*, 46: 126-40.
- Papineau, David [1993], *Philosophical Naturalism*, Oxford: Blackwell.
- Parfit, Derek [1984], *Reasons and Persons*, Oxford: Oxford University Press.
- [2011], *On What Matters*, Vol. 1, Oxford: Oxford University Press.
- Pritchard, Duncan [2015], ‘Risk’, *Metaphilosophy*, 46: 436-61.
- Quine, W. V. [1960], *Word and Object*, Cambridge, MA: MIT Press.
- Raz, Joseph [1986], ‘Value Incommensurability: Some Preliminaries’, *Proceedings of the Aristotelian Society*, 86: 117-34.
- [1997], ‘Incommensurability and Agency’, in Ruth Chang (ed.), *Incommensurability, Incomparability, and Practical Reasoning*, Cambridge, MA: Harvard University Press: 110-28.
- [1999], *Engaging Reason*, Oxford: Oxford University Press.
- Regan, Donald [1997], ‘Value, Comparability, and Choice’, in Ruth Chang (ed.), *Incommensurability, Incomparability, and Practical Reasoning*, Cambridge, MA: Harvard University Press: 129-50.
- Regan, Tom [1983], *The Case for Animal Rights*, Berkeley: University of California Press.
- Richter, Hans [1965] (2016), *Dada: Art and Anti-Art*, London: Thames & Hudson.
- Rinard, Susanna [2015], ‘A Decision Theory for Imprecise Probabilities’, *Philosophers’ Imprint*, 15: 1-16.
- Rose, James et al. [2014], ‘Can Fish Really Feel Pain?’, *Fish and Fisheries*, 15: 97-133.
- Rosen, Gideon, and Smith, Nicholas [2004], ‘Wordly Indeterminacy: A Rough Guide’, *Australasian Journal of Philosophy*, 82: 185-98.
- Sainsbury, Mark [1997], ‘Easy Possibilities’, *Philosophy and Phenomenological Research*, 57: 907-19.
- Schärp, Kevin [2013], *Replacing Truth*, Oxford: Oxford University Press.
- Schiffer, Stephen [2003], *The Things We Mean*, Oxford: Oxford University Press.
- Schoenfield, Miriam [2016], ‘Moral Vagueness is Ontic Vagueness’, *Ethics*, 126: 257-82.
- Searle, John [1992], *The Rediscovery of the Mind*, Cambridge, MA: MIT Press.
- Sen, Amartya [1997], ‘Maximization and the Act of Choice’, *Econometrica*, 65: 745-79.
- Shelley, James [2003], ‘The Problem of Non-Perceptual Art’, *British Journal of Aesthetics*, 43: 363-78.
- Shpall, Sam [2014], ‘Moral and Rational Commitment’, *Philosophy and Phenomenological Research*, 88: 146-72.
- Sidgwick, Henry [1962] (1907), *The Methods of Ethics*, 7<sup>th</sup> ed., London: Palgrave Macmillan.
- Simon, Jonathan [2017], ‘Vagueness and Zombies: Why ‘Phenomenally Conscious’ has no Borderline Cases’, *Philosophical Studies*, 174: 2105-23.

- Singer, Peter [2002], *Animal Liberation*, 3<sup>rd</sup> ed., New York: Harper Collins.
- [2011], *Practical Ethics*, 3<sup>rd</sup> ed., Cambridge: Cambridge University Press.
- Sinnott-Armstrong, Walter [1988], *Moral Dilemmas*, Oxford: Blackwell.
- Sorensen, Roy [2001], *Vagueness and Contradiction*, Oxford: Oxford University Press.
- Stecker, Robert [1997], *Artworks: Definition, Meaning, Value*, University Park: Pennsylvania State University Press.
- Steinbock, Bonnie [2011], *Life Before Death: The Moral and Legal Status of Embryos and Fetuses*, 2<sup>nd</sup> ed., New York: Oxford University Press.
- Sud, Rohan [2019], ‘Moral Vagueness as Semantic Vagueness’, *Ethics*, 129: 684-705.
- Sumner, L. W. [1983], *Abortion and Moral Theory*, Princeton: Princeton University Press.
- [1985], ‘Moral Theory and Moral Standing: A Reply to Woods and Soles’, *Dialogue*, 24: 691-9.
- Sutton, Jonathan [2005], ‘Stick to What You Know’, *Noûs*, 39: 359-96.
- Taylor, David [2018], ‘A Minimal Characterization of Indeterminacy’, *Philosophers’ Imprint*, 18: 1-25.
- Thomson, Judith Jarvis [1971], ‘A Defense of Abortion’, *Philosophy & Public Affairs*, 1: 47-66.
- Tooley, Michael [1972], ‘Abortion and Infanticide’, *Philosophy & Public Affairs*, 2: 37-65.
- Tye, Michael [1990], ‘Vague Objects’, *Mind*, 99: 535-57.
- [1996], ‘Is Consciousness Vague or Arbitrary?’, *Philosophy and Phenomenological Research*, 56: 679-85.
- [2000], *Consciousness, Color and Content*, Cambridge, MA: MIT Press.
- [2017], *Tensed Bees and Shell-Shocked Crabs*, New York: Oxford University Press.
- Van Inwagen, Peter [1990], *Material Beings*, Ithaca, NY: Cornell University Press.
- Vasari, Giorgio [1550] (1991), *The Lives of the Artists*, trans. Julia Conaway Bondanella and Peter Bondanella, New York: Oxford University Press.
- Warren, Mary Anne [1973], ‘On the Moral and Legal Status of Abortion’, *The Monist*, 57: 43-61.
- Williams, Robert [2008], ‘Multiple Actualities and Ontically Vague Identity’, *Philosophical Quarterly*, 58: 134-54.
- [2012], ‘Indeterminacy and Normative Silence’, *Analysis*, 72: 217-25.
- [2014a], ‘Decision-Making under Indeterminacy’, *Philosophers’ Imprint*, 14: 1-34.
- [2014b], ‘Nonclassical Minds and Indeterminate Survival’, *Philosophical Review*, 123: 379-428.
- [2016], ‘Angst, Indeterminacy and Conflicting Values’, *Ratio*, 29: 412-33.
- [2017], ‘Indeterminate Oughts’, *Ethics*, 127: 645-73.
- Williamson, Timothy [1994], *Vagueness*, London: Routledge.
- [2000], *Knowledge and Its Limits*, New York: Oxford University Press.
- [2004], ‘Reply to McGee and McLaughlin’, *Linguistics and Philosophy*, 27: 113-22.
- [2007], ‘Knowledge Within the Margin for Error’, *Mind*, 116: 723-26.
- [2009], ‘Probability and Danger’, *The Amherst Lecture in Philosophy* 4: 1-35  
[<amherstlecture.org/williamson2009/>](http://amherstlecture.org/williamson2009/).
- Wood, Beatrice [1985], *I Shock Myself*, Ojai, CA: Dillingham Press.
- Wright, Crispin [2001], ‘On Being in a Quandary’, *Mind*, 110: 48-98.
- Zangwill, Nick [2007], *Aesthetic Creation*, Oxford: Oxford University Press.
- Zimmerman, Michael [2002], ‘Taking Luck Seriously’, *Journal of Philosophy*, 99: 553-76.