Self-Referential Probability and Rationality

Miriam Shamsy Jane Bowen

Submitted in accordance with the requirements for the degree of Doctor of Philosophy

The University of Leeds

School of Philosophy, Religion and History of Science

October 2021

The candidate confirms that the work submitted is her own and that appropriate credit has been given where reference has been made to the work of others.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

The right of Miriam Shamsy Jane Bowen to be identified as Author of this work has been asserted by Miriam Shamsy Jane Bowen in accordance with the Copyright, Designs and Patents Act 1988.

To my mother

Acknowledgements

I was fortunate enough to be able to draw on three sources of guidance for this thesis: Robbie Williams, Ed Elliott (who took over for Jack in my second year) and Jack Woods. I'm very grateful to Robbie who has been enthusiastic about the project from the start and has helped guide me and shape my approach to philosophy over the last four years. This thesis wouldn't have been possible without him. Many thanks also to Ed. Our discussions and his comments have been central to developing the ideas in this thesis. Thanks to Jack for his help focusing the direction of this project in my first year. Many thanks also to my examiners Anna Mahtani and Darren Bradley.

My work was made possible by The University of Leeds Research Scholarship

I've been lucky to meet so many wonderful people at Leeds and had such a supportive graduate research environment. There are many people deserving of mention, these include: Adina Covaci, Andreas Bruns, Alice Murphy, Nathan Wood, Aleksander Domoslawski, Jyl Schacher, Kevin Jones, Will Gamester, Sam Mason, Ludovica Adamo and Thomas Brouwer. Particular thanks to David Heering, Arthur Carlyle and Alex Stamatiadis-Bréhier who have all been excellent friends and helped me stay sane through the challenges of completing a PhD in a pandemic.

I started to gain interest in theories of truth and philosophy of probability while doing my masters in Bristol. I'd like to thank Catrin Campbell-

Moore for introducing me to self-referential probability and our discussions since.

My thanks to my school friends and Scarlett Lawson in particular who has planned trips, organised gatherings and is always up for a day exploring North Wales when we're both back home. Many thanks also to Rochyne Delaney McNulty who has also been an excellent friend, housemate and climbing partner.

A big thank you to my Warwick Climbing friends. Our climbing trips are always such a joy and much needed time away from looking at a laptop.

Thanks to my brother Abbas Bowen for all his help and support, and finally I would like to thank my mother, Soussan Khododadeh, who has helped me more than I can say. She has always encouraged me in everything I do.

Abstract

This thesis argues for attitudes an agent ought to adopt towards two problematic examples of probabilistic self-reference. In particular, I look at a case of self-referential probability I refer to as the Probabilistic Liar, due to its similarities to the Liar paradox. The Probabilistic Liar emerges when an agent's credence can act as evidence for the truth of the proposition. Examples of self-reference turn out to be problematic for traditional Bayesian accounts of rationality.

I develop an account of how rational agents ought to respond to the Probabilistic Liar by *suspending judgment*. Suspended judgment is an attitude more naturally talked about in traditional all-or-nothing belief models. I argue for suspended judgment in a credal framework and in particular that suspended judgment is a determinate attitude that should be represented by imprecise credences. This gives a principled way of weakening the requirement that a rational agent's degrees of beliefs ought to be probabilistically coherent.

Once a solution to the Probabilistic Liar has been given a new question emerges. Can we give another example of problematic probabilistic self-reference in terms of the suspended judgement attitude? That is, can we give a Revenge problem. I explore how a Revenge problem can be generated for my account and how a rational agent can respond by having *indeterminate attitudes*.

Finally, I argue that both the Probabilistic Liar and Revenge problems are cases of indeterminacy. I then look at the normative question of what attitude an agent ought to adopt towards cases of indeterminacy. Drawing on the attitudes I have argued for in the thesis, I argue for a *pluralist* answer to the normative question.

Contents

General Introduction **1**

1 Suspended Judgment **11**

1.1 Introduction 11

1.2 Probabilistic Liar 14

1.2.1 Consequence of self-referential probability 15

1.2.2 Parallels to the Liar Paradox 24

1.3 Suspended Judgment 26

1.3.1 Norms for suspended judgment 29

1.3.2 Coherence Principles in a Belief Framework 32

1.4 Suspended Judgment and Degrees of Belief? 34

1.4.1 Coherence Principles in a Degree of Belief Framework 37

1.5 Suspension as Imprecise Credence 40

1.5.1 Why Imprecise Credences ? 43

1.5.2 Imprecise credences and suspended judgment . . 45

1.6 Suspending Judgment in the Probabilistic Liar 59

1.7 Conclusion 65

2 Imprecise Credences and Comparativism **67**

2.1 Introduction 67

2.2 Intersectionism 69

2.3 Degrees of Belief 71

2.3.1 What are degrees of belief 72

2.3.2 Desiderata for an account 78
2.4 Comparativism 83
2.4.1 Comparative Probability 85
2.4.2 Ordinal vs Cardinal 88
2.5 Problems for Comparativism 93
2.5.1 Multiple Probabilistic Representations 94
2.5.2 No Probabilistic Representation 100
2.5.3 Interpersonal Comparability 106
2.6 Conclusion 121

3 Intersectionism and Suspended Judgment 123

3.1 Introduction 123
3.2 Imprecise Credences 125
3.3 Supervaluationism 128
3.4 Intersectionism 132
3.4.1 Problems for intersectionism 138
3.5 Suspended Judgment 139
3.6 Conclusion 143

4 Revenge Paradox 145

4.1 Introduction 145
4.2 Revenge for the Liar 147
4.2.1 Solutions to Revenge for the Liar 151
4.3 Revenge for the Probabilistic Liar 159
4.4 Possible Solutions to Revenge for the Probabilistic Liar . . 163
4.5 A solution to Revenge for the Probabilistic Liar 166
4.5.1 Supervaluationism 168
4.6 Revenge in terms of indeterminacy 174

4.6.1 Higher-order Indeterminacy 180
4.6.2 The logic of determinacy 184
4.7 Indeterminacy 193

5 The Cognitive Role of Indeterminacy **195**

5.1 Introduction 195
5.2 Probabilistic Liar, Revenge and Indeterminacy 203
5.2.1 Indeterminacy 203
5.2.2 Probabilistic Liar 207
5.2.3 Revenge 210
5.3 Monism and Pluralism 213
5.3.1 Monism 215
5.4 Pluralism 220
5.4.1 Examples of Indeterminacy 221
5.4.2 Domain Specific Pluralism 224
5.4.3 Interaction Effects 229
5.5 Modest Pluralism 235
5.5.1 Modest Pluralism Applied 238
5.6 Probabilistic Liar, Revenge and Indeterminacy again 240
5.6.1 Same attitude 241
5.6.2 Not both indeterminate 242
5.7 Conclusion 243

General Conclusion **246**

Appendix **252**

References **257**

General Introduction

The following are examples of self-reference:

(1) What sentence (1) says is not true.

(2) What sentence (2) says is true.

These are examples of self-reference in truth. Sentence (1) expresses a version of the Liar Paradox, any truth value assigned to this swiftly leads to contradiction. (2) expresses the Truth Teller sentence where any specific assignment of truth-status is apt to seem unprincipled. There are also analogous examples of self-reference that involve belief:

(3) I believe what sentence (3) says.

(4) I do not believe what sentence (4) says.

These are examples of doxastic self-reference that involve flat out belief states.[1] My main focus here will be doxastic self-reference that involves degrees of belief: probabilistic self-reference.

Probabilistic self-reference and the above examples of doxastic self-reference occur when the truth or the chance of an event is dependent on the degree of belief (credence) / belief of the agent. Greaves describes these as situations where the agent is not a *pure observer*.[2]

[1] That is to say, in traditional accounts of belief where an agent can either believe or disbelieve.

[2] These are not the only scenarios where probabilistic self-reference arises. We can generate probabilistic self-reference in languages in the following way:

We can give examples of 'real world' situations where it appears that an agent's credence changes the likelihood of an event. There are cases that look like probabilistic truth-tellers:

> *Leap*:
> James is stuck in the Alps and the only way he can escape is to jump across a chasm. If James is confident that he will make the leap he will attempt a proper run up and thus it makes it more likely that he will actually make the leap. If James is unconfident he will make the leap then he will not commit, he might stumble on his run and this makes it less likely that he will make the leap. We can be more precise and say that whatever the degree of confidence $x \in [0, 1]$ James has in making the leap he has chance x of making the leap.[3]

One might think in this situation that there are only two rational beliefs James could have. If he has credence 1 then he definitely makes the leap and if he has credence 0 he definitely does not make the leap. When he has the extremal credences his credence matches the truth.[4] However, one could also argue that *any* credence is rational, not just the extremal credences since any credence matches the

(π) The probability of the sentence (π) is less than 0.5.

In this thesis I will focus on self-reference that occurs in Greaves-style scenarios.

[3] An adaptation of an example from James (1956) and (Greaves, 2013, p.916).

[4] As I will discuss in Chapter 1, section 1.2 one might want to maximise the accuracy of their beliefs, therefore making the extremal credences the only possible rational attitudes to adopt.

chance.[5] Or, one could argue that the only rational credence James can have is credence 1 since this is the only case where it is guaranteed he will successfully make the leap.[6] It is therefore not clear what credences it is rational for James to have.

We can also give real world examples where there is an inverse relationship between the chance of the event and the agent's credence. Once again, turning to Greaves, she gives the following example.[7]

> *Promotion*:
> Alice is up for promotion. Her boss, however, is a deeply insecure type: he is more likely to promote Alice if she comes across as lacking in confidence. Furthermore, Alice is useless at play-acting, so she will come across that way iff she really does have a low degree of belief that she's going to get the promotion. Specifically, the chance of her getting the promotion will be $(1 - x)$ where x is whatever degree of belief she chooses to have in the proposition P that she will be promoted. What credence in P is it epistemically ratio-

[5] This could be argued for by appeal to a chance-credence principle, I will discuss chance credence principles in more detail in Chapter 1, section 1.4.

[6] We could think of this as epistemic consequentialism where adopting a belief is rational in so far as it brings about *epistemic value*.

> ... a state of affairs is one of high epistemic value for a given agent just in case it is a state of affairs in which there is a good degree of fit between that agent's beliefs and the truth. (Greaves, 2013, p.919)

[7] Egan and Elga (2005) also give an example about an agent AE who is anti-reliable about direction. AE's all-things-considered judgment about direction is wrong, so when he gets to an intersection and needs to make a decision about what direction to turn any beliefs he has about what direction to turn will lead him astray.

nal for Alice to have? (Greaves, 2013, pp.915-916)

Assuming that Alice is aware of the setup of the situation and aware of her own degrees of belief she can take her degree of belief as evidence for the chance of her getting the promotion. In this example it seems like there is a unique epistemically rational credence for Alice to have, $x = 0.5$. Only when she has degree of belief 0.5 that she will get the promotion does her belief match the chance of her getting the promotion.

If we consider an agent's beliefs in degrees then we would like to be able to express self-reference so we can express scenarios such as *Leap* and *Promotion*. Probabilistic self-reference can however be problematic, and we can give examples where there is no clear answer to what a rational agent ought to believe. A particularly problematic probabilistic self-referential scenario is that of the Probabilistic Liar, a probabilistic analogue of the Liar paradox.

In this thesis I focus on a version of the Probabilistic Liar that arises when an agent stipulates that a proposition (α) is true if and only if they have a low credence in (α).

$$(\alpha) \quad Cr(\alpha) < 0.5$$

It is unclear what attitude an agent ought to adopt in (α). If the agent adopts a credence of less than 0.5 in (α) then (α) is true and so it seems she ought to have in fact adopted a high credence in (α). On the other hand if she adopts a credence greater than or equal to 0.5 in (α) this makes (α) false and so it seems she ought to in fact have a low credence in (α). For any precise credence she adopts it seems like she ought to change her credence to reflect that she is certain that (α) is true or false. It seems there is no stable attitude she can adopt towards (α). This seems puzzling, since the agent knows the setup of the proposition in advance of adopting a credence towards it.

The source of our puzzlement over what attitude to adopt in (α) might arise from the problems it causes between two plausible principles of rationality: Rational Introspection and Probabilism. Caie (2013) shows that an agent cannot have good introspective access to her own beliefs and have probabilistically coherent beliefs in the Probabilistic Liar.

In Chapter 1 I present this puzzle and argue for a solution - that an agent ought to suspend judgment in (α). Suspended Judgement is more clearly understood in a traditional belief framework where an agent either believes, disbelieves, or suspends judgement. I give an account of how suspended judgment can be understood in a credal framework. In particular, I argue that it should be represented by imprecise credences. Imprecise credences represent an agent's attitudes with a set of credences rather than a precise credences function. This also provides a specific weakening of the norm of Probabilism.

In chapters 2 and 3 I turn in more detail to the underlying interpretation of imprecise credences. There are a number of accounts of imprecise credences. In order to give a full account of the suspended judgment attitude a particular interpretation of imprecise credences must be given and justified. I argue for a particular interpretation:

Comparativist Intersectionism. A key feature of this interpretation is that the credal set as a whole is taken to represent the agent's determinate attitude.

The comparativist intersectionist interpretation looks at the comparative information contained in the credal set. There are two assumptions made in the comparativist intersectionist interpretation:

(i) *Comparativism*: agent's have comparative beliefs (which can explain their partial beliefs).[8]

(ii) These comparative beliefs may be incomplete due to an agent lacking certain comparative beliefs (and we therefore need to be able to represent that there are gaps in the comparative belief ordering).

Chapter 2 focuses upon the first assumption of the comparativist intersectionist interpretation. It therefore focuses on the foundational question of what degrees of belief are. I argue that *comparativism* is a viable account of what degrees of belief are.

In Chapter 3 I turn to the second part of justifying comparativist intersectionism. If one accepts comparativism then a desirable feature of an account is that it accommodates the possibility that agents may not have complete comparative belief orderings. Precise credal representations cannot accommodate this, so this suggests turning to imprecise credences. Comparativism therefore motivates having an

[8]Comparative beliefs are beliefs of the form: 'I think *A* is at least as likely as *B*'. Partial beliefs are beliefs of the form: 'I am very confident that *A*' or 'I am slightly confident that *B*'.

imprecise credal set. There are however multiple interpretations of the imprecise credal set that are compatible with an agent not having a complete comparative belief ordering.

The supervaluationist interpretation takes the set of credences as representing that an agent has indeterminate attitudes. This interpretation treats each credence function in the credal set as a permissible precisification of the agent's attitude. When the credence functions in the set disagree on what credence they assign to a particular proposition we interpret this as representing that the agent has a vague or indeterminate attitude towards that proposition. One explanation we can give for an agent having an incomplete comparative belief ordering is that her comparative beliefs are sometimes indeterminate.

Another interpretation of the imprecise credal set is the intersectionist interpretation. This interpretation holds that an agent may determinately lack certain comparative beliefs. On this interpretation it is not simply the case that it is vague whether a certain comparative ordering holds between two propositions. Rather, it represents that an agent determinately lacks certain comparative beliefs. I argue that an intersectionist interpretation naturally follows from a comparativist view of belief.

In Chapter 4 I turn to the question of whether a Revenge problem can be generated for my account. Revenge problems are familiar from the Liar paradox: once a solution has been offered to the Liar paradox, we can ask whether a new and strengthened Liar sentence can be given which applies to / causes problems for that solution. Given the solution of Suspended Judgment, can a strengthened Probabilistic Liar in terms of suspended judgment be given?

It looks like it will be simple to generate a Revenge problem in terms of suspended judgment. We simply generate a new proposition with probabilistic self-reference, but this time referring to the content of the agent's credal set. The solution of suspending judgment will not act as a solution to this. I argue instead that an agent ought to have an indeterminate attitude towards it. Indeterminacy in an agent's attitude can be represented by imprecise credences with a supervaluationist interpretation of a set of credal sets.

Once we have given a solution to Revenge for the Probabilistic Liar we can ask whether a newly strengthened Revenge problem can be given in terms of having indeterminate attitudes. I argue that we can iterate the indeterminacy to generate a series of Revenge problems. Each of these new Revenge problems requires a solution, and I argue that the same iteration can be applied to my proposed solution. In the original Revenge problem the solution was that the agent has indeterminate attitudes about the content of the credal set. For each subsequent Revenge problem the solution again is that the agent has an indeterminate attitude to the proposition in question.

When we consider the Probabilistic Liar and Revenge problems they both seem like cases of indeterminacy. In Chapter 5 I argue that indeterminacy should be understood as an umbrella term that encompasses a range of related phenomena. If they are both the same type of phenomena this raises the question of whether there is a norm that guides what attitude an agent ought to adopt towards cases of indeterminacy. If there is, should we expect there to be a unique attitude an agent ought to adopt in virtue of being certain a proposition is indeterminate? Following Williams (2012) I refer to this as the cognitive role of indeterminacy.

In the previous chapters I argue for two distinct attitudes: suspended judgment and having indeterminate attitudes.[9] If there is a unique attitude that an agent ought to adopt towards indeterminacy then this looks problematic for my account. In this chapter I argue for a pluralist stance on the cognitive role of indeterminacy. This position leaves open the possibility that there may be multiple norms that govern what attitude an agent ought to adopt towards cases of indeterminacy.

My focus in this thesis is the Probabilistic Liar. There are also strong parallels between the Probabilistic Liar and the Liar paradox which I will discuss in more detail in Chapter 1. Given this, part of my methodology is to take inspiration from the Liar. However, notably I am not trying to solve them together. In light of the parallels between the Liar and the Probabilistic Liar there is reason to think they should have similar solutions. Different solutions are not to be ruled out however and the question of how my proposed solution to the Probabilistic Liar and different Liar solutions interact is beyond the scope of this

[9] One question that arises here is, since I argue an agent ought to adopt an indeterminate attitude in the Revenge cases, then why not in the original case too? I address this question at the end of Chapter 5.

thesis.[10]

In Chapters 1 and 4 I look to styles of solution for the Liar paradox and Revenge for the Liar. As I will discuss in Chapter 1 section 1.2.2 I am broadly sympathetic to failure of excluded middle in the Liar (such as the account given by Maudlin (2004)). However, I will be sticking to classical logic for the credal case.[11] As I will discuss in Chapter 4 section 4.2.1 I think a hierarchy approach is a good solution for Revenge for the Liar (such as the account offered by Cook (2007)). In Revenge for the Probabilistic Liar, I think we can take inspiration form this and give a solution to Revenge for the Probabilistic Liar by taking iterations of indeterminacy. While the Liar helps provide inspiration this thesis focuses on providing an account of a solution to the Probabilistic Liar independently of a defence of any particular solution to the Liar paradox.

[10]There are also important disanalogies between the Liar and the Probabilistic Liar such as the credal Probabilistic Liar being connected to an agent's psychology and the possible attitudes agents can have - I give some discussion to this in Chapter 4 section 4.4. Considerations such as this might give us reasons for thinking they merit different solutions. I also discuss in Chapter 4 how there is disunity between my response towards the Probabilistic Liar and Revenge for the Probabilistic Liar so my account for the Probabilistic Liar alone is not uniform.

[11]If I do commit to Maudlin's non-classical approach this opens up some substantial questions, namely the question of combining a non-classical logic for the liar with classical probabilism.

Chapter 1
Suspended Judgment

1.1 Introduction

In the General Introduction I gave some 'real world' examples of self-reference probability from Greaves (2013). These examples can be unproblematic. We can see there is some attitude that the agent ought to adopt in the scenario. The *Promotion* example I gave looks like an unproblematic example of self-reference where there is an inverse relation between the chance of an event and the agent's credence:

> *Promotion*:
> Alice is up for promotion. Her boss, however, is a deeply insecure type: he is more likely to promote Alice if she comes across as lacking in confidence. Furthermore, Alice is useless at play-acting, so she will come across that way iff she really does have a low degree of belief that she's going to get the promotion. Specifically, the chance of her getting the promotion will be $(1 - x)$ where x is whatever degree of belief she chooses to have in the proposition P that she will be promoted. What credence in P is it epistemically rational for Alice to have? (Greaves, 2013, pp.915-916)

We can amend this example to give a more problematic scenario:[1]

> *Promotion**:
> Alice is up for promotion. Her boss, however, is a deeply insecure type: he is more likely to promote Alice if she comes across as lacking in confidence. Furthermore, Alice is useless at play-acting, so she will come across that way iff she really does have a low degree of belief that she's going to get the promotion. Specifically, she will get the promotion if her credence in getting the promotion is less than 0.5 and she will not get the promotion if her credence in getting the promotion is greater than or equal to 0.5.

In this case there is no clear answer about what attitude Alice should adopt. We can give a formal version of this which I will refer to as the Probabilistic Liar, which I present in section 1.2. I will outline the intuitive problem that arises from the Probabilistic Liar and also show that it gives rise to contradictions with the plausible norms of rationality; Probabilism and Rational Introspection.

In section 1.2.1 I outline Caie's (2013) presentation of the conflict between Probabilism and Rational Introspection. Caie suggests weakening rationality requirements and accepting that in some cases a rational agent can have probabilistically incoherent beliefs. His account therefore rejects Probabilism in a very strong way in order to accommodate self-referential probability. I show there is an alternative way of accommodating self-referential probability that also rejects Prob-

[1]A similar example can be found in (Konek and Levinstein, 2017, p.37).

abilism but in a much weaker way than Caie proposes and in a way that retains the *essence* of probabilism. To do this I look to alethic self-reference and consider the parallels between the Liar paradox and the Probabilistic Liar in more detail in section 1.2.2. In particular, one approach to alethic self-reference is to consider three-valued valuation schemes. This suggests a possible way of addressing the Probabilistic Liar, to consider a doxastic attitude that is not a degree of belief: suspended judgment.

In section 1.3 I present examples that illustrate when suspension of judgment is called for and its role in a traditional (all-or-nothing) belief framework. In sections 1.3.1 and 1.3.2 I provide norms and coherence conditions for suspended judgment in light of these. We can also consider how to understand suspended judgment in a credal model. In section 1.4 I consider how we can express suspended judgement in credal terms. I argue that in the context of degrees of belief we should not equate suspended judgment as a middling credence. However, we can still understand suspended judgment in a credal context by treating suspended judgment as having an imprecise credence. In section 1.5 I outline a general account of what imprecise credences are and in section 1.5.1 some independent reasons for why we should adopt imprecise credences. Using this in section 1.5.2 I show how suspended judgment can be understood in terms of imprecise credences. In light of this understanding in section 1.6 I return to the Probabilistic Liar example showing how suspended judgment acts as a solution.

1.2 Probabilistic Liar

In this thesis I will focus on a Probabilistic Liar example that arises when the truth of a sentence is equivalent to an agent having a low credence in that sentence. We can give a formal version of this as follows (an adaptation of an example by Caie (2013)).

We imagine an agent Alex who stipulates the meaning of a sentence (α) to be the following:

> (α) Alex's credence in the proposition expressed by (α) is less than 0.5

Abbreviating 'Alex's credence that ...' to '*Cr*' and 'the proposition expressed by' to 'ρ' we can write this more concisely as:

> (α) $Cr\rho(\alpha) < 0.5$.

What credence should Alex have in (α)?[2]

If Alex's credence in (α) is less than 0.5 it follows that (α) is true (since Alex's credence in the proposition expressed by (α) *is* less than 0.5).

---

[2] I am making the assumption that (α) expresses a proposition. One might think that there are problems with this presupposition given the self-referential nature of (α). Caie argues that:

> A sufficient condition for [a sentence] (ϕ) to express a proposition is if (ϕ) can be embedded under metaphysical modal operators in a way that results in a true sentence. (Caie, 2013, p.540)

He goes on to give examples and argue that self-referential sentences are "not barred from expressing a proposition simply in virtue of containing a term that purports to refer to the proposition expressed by that sentence" (ibid, p.541). See also (Campbell-Moore, 2015b, p.395 fn.4) where she notes that even if one disagrees with Caie's assumption that (α) expresses a proposition we can get a similar result by instead viewing Alex as assigning credences to sentences of a language that has predicates like '$Cr(\cdot) < 0.5$' that apply to codes of sentences (taking a background theory of arithmetic coding of sentences).

Having a credence less than 0.5 is therefore problematic since we do not want to recommend that Alex have a 'low' credence in a sentence she knows to be true.[3] Indeed, in light of the above reasoning and the fact that we know (α) is now true it seems reasonable to try assigning credence 1 to (α).[4] However, when Alex has credence 1 in (α), or indeed if she has any credence greater than or equal to 0.5 in (α) it follows that (α) is false. If (α) is false then Alex should have credence 0 in (α).[5] We have already seen that if she has credence less than 0.5 this makes (α) true. This reasoning leads to a constant flipping between credences and truth values and we can see that any degree of belief that Alex settles on will be subject to this instability. For sentences such as (α) we can see that there is a difficulty recommending a particular degree of belief for an agent to have, indeed there does not appear to be *any* stable credence we can recommend Alex to have.

1.2.1 Consequence of self-referential probability

We can clearly see that there is a difficulty in answering the question 'What credence should Alex have in (α)?'. There are also broader problems that self-referential sentences pose for theories of rationality. In particular, Caie shows we can sharpen the intuition that there is something problematic with assigning any precise credence to (α) by showing that in a Probabilistic Liar example there is a conflict be-

[3] By low credence I mean a credence less than 0.5.

[4] One could make the weaker claim that Alex ought to be confident in (α) (have a credence greater than 0.5) if she is certain (α) is true. Even this weaker claim proves problematic.

[5] Or indeed, one could again make the weaker claim that Alex ought to have low credence in (α) and this is also problematic.

tween two plausible norms of rationality.[6]

Caie's example goes as follows, we consider an agent Hiro who stipulates the meaning of the sentence (#):

> (#) Hiro's credence in the proposition expressed by (#) isn't greater than or equal to 0.5

or

> (#) $\neg Cr\rho(\#) \geq 0.5$

there is a conflict between two highly plausible norms of rationality; Probabilism and Rational Introspection.

Probabilism is the claim that a rational agent's credences form a probability function, by which I mean the credence function $Cr(\cdot)$ satisfies the following for propositions in the domain of the credence function:

- Normalisation: For $\top$ a logical truth, $Cr(\top) = 1$
- Non-negativity: The probability of any proposition ϕ is a non-negative real number.
- Finite additivity: If ϕ and ψ are incompatible propositions, then $Cr(\phi \vee \psi) = Cr(\phi) + Cr(\psi)$.

Having probabilistically coherent beliefs seems like a reasonable constraint on rational belief. Indeed, it seems strange to endorse the view that a rational agent could have credence 0.9 in a proposition P and also have a credence greater than 0.1 in $\neg P$. There have also been a number of arguments offered in favour of Probabilism such as Dutch

[6] This example is a strengthened version of the Alex example I presented above.

Book arguments and Accuracy Dominance arguments.

Dutch Book arguments show that if an agent's credences do not obey Probabilism then there are a set of bets that an agent can be offered (and take at her fair betting price) that guarantees her a sure loss.[7] Accuracy dominance arguments assume that accuracy is an epistemic virtue (Pettigrew (2016) describes accuracy as the "fundamental epistemic virtue" (2016, p.6)). An agent is irrational if they have a credence Cr which is accuracy dominated.[8] The core idea of accuracy arguments is that by giving an account of what it means for degrees of belief to accurately represent the world we can explain why having probabilistically coherent beliefs "contributes to the basic epistemic goal of accuracy" (Joyce, 1998, p.567).

Rational Introspection is the highly plausible claim that "a rational agent must be responsive to her own credal state" (Caie, 2013, p.530). This means a rational agent should be able to know or be confident what credence they currently assign to a certain proposition. I will refer to Rational Introspection as a norm, however the arguments I will present that involve Rational Introspection do not require that it is taken as a norm of rationality. It need not be the case that a rational agent *must* have this kind of introspective capacity. Rather, all that is really needed is that it's *possible* for an ideally rational agent to have this introspective capacity.

Caie shows that with self-referential situations like the Hiro example

[7]See Vineberg (2016) and (Titelbaum, 2019, §3.2).

[8]See Joyce (1998) and Pettigrew (2016).

there is a conflict between Rational Introspection and Probabilism. He uses this as a starting point to argue that we should drop the condition of Probabilism and instead accept rational probabilistic incoherence.

The argument goes as follows. We consider the following:

$$(\#)\ \ \neg Cr\rho(\#) \geq 0.5$$

Where 'ρ' is understood as 'the proposition expressed by'. So we have:

$$(1)\ \ \rho(\#) = \rho\text{`}\neg Cr\rho(\#) \geq 0.5\text{'}$$

Caie uses the following positive and negative sensitivity introspection principles:

$$(2)\ \ [Cr\rho(\#) \geq 0.5] \rightarrow [Cr(\rho\ \text{`}Cr\rho(\#) \geq 0.5\text{'}) > 0.5]$$

$$(3)\ \ [\neg Cr\rho(\#) \geq 0.5] \rightarrow [Cr(\rho\ \text{`}\neg Cr\rho(\#) \geq 0.5\text{'}) > 0.5]$$

Where there is an assumption in (3) that the relevant propositions are all in the domain of the credence function.[9] Each of these introspection principles looks plausible. If an agent has $Cr\rho(\#) \geq 0.5$ then when they introspect they should be confident that they have $Cr\rho(\#) \geq 0.5$ (and similarly in the case of negative introspection).[10]

[9]When this is the case the negative introspection principle looks like a reasonable assumption. If we do not include this assumption then there is the possibility that $\neg Cr\rho(\#) \geq 0.5$ because $Cr\rho(\#) \geq 0.5$ is not defined. In this case, the introspection seems much less plausible.

[10]Note Caie's introspection principles can be viewed in two ways. We could think of

Assume

$$(4)\ \neg Cr\rho(\#) \geq 0.5$$

Then by (1) $\neg Cr(\rho`\neg Cr\rho(\#) \geq 0.5') \geq 0.5$

but given our assumption of (3) we get a contradiction.

So it follows that

$$(5)\ Cr\rho(\#) \geq 0.5$$

by (5) and (1) it follows that $Cr(\rho`\neg Cr\rho(\#) \geq 0.5') \geq 0.5$

and by (5) and (2) it follows that $Cr(\rho`Cr\rho(\#) \geq 0.5') > 0.5$

so $Cr(\rho`\neg Cr\rho(\#) \geq 0.5') + Cr(\rho`Cr\rho(\#) \geq 0.5') > 1$

So, Hiro's beliefs are not probabilistically coherent. This shows that for Hiro to have probabilistically coherent beliefs about (#) there must be some limitation on his introspective abilities. Caie takes this conflict as evidence that we should reject Probabilism rather than reject Rational Introspection. If Rational Introspection fails it is not just that an agent has poor access to their own beliefs, it means that in some situations an agent is rationally required to have poor access to their own beliefs. Caie claims that it is prima facie implausible that this is a requirement on an agent's rationality. This seems right, it ought to be *possible* for a rational agent to introspect on their own determinate

this as expressing how a real agent updates their attitudes when they introspect. In this case the process takes place over time, the agent has an attitude, they introspect on that attitude and then come to have an attitude about their own beliefs based on their introspection. We could also view it as expressing a material conditional in the case of an ideal agent. That is, if we assume that an ideal agent has immediate introspective access to their epistemic state and updates instantaneously.

attitudes.

While the conflict that Caie presents shows that we must weaken one of Rational Introspection or Probabilism I have a conflicting intuition to him regarding how we weaken one of these norms as it also appears implausible that an agent has sharp credences and be rationally required to accept $Cr(\neg p) \geq 0.5$ and $Cr(p) > 0.5$.

Caie gives a more worrying argument against Probabilism by arguing that Accuracy Dominance arguments—which have been traditionally used to argue *for* Probabilism—in fact, show that the most accurate credence Hiro can have is a probabilistically incoherent credence. He shows this by considering the smallest algebra that contains (#), $\top$ and $\bot$. We assume that Hiro has credence 1 in $\top$ and credence 0 in $\bot$. We can consider the following possible worlds: w_1 where the proposition expressed by (#) is true and w_2 where the proposition expressed by (#) is false. Thinking of this graphically, we can take the x-coordinate to represent the negation of the proposition expressed by (#) and the y-coordinate to be the proposition expressed by (#) then, $w_1 = \langle 0, 1 \rangle$ and $w_2 = \langle 1, 0 \rangle$. The most accurate belief Hiro can have with regards to (#) and ¬(#) is $b = (1, 0.5)$ which is a probabilistically incoherent belief.[11] There are a number of ways we might push back against Caie's conclusion that accuracy-dominance suggests we should be probabilistically incoherent. One possibility is to reject the particular way Caie applies accuracy considerations to self-referential sentences. Campbell-Moore (2015b) also presents argu-

[11] Caie is assuming that such a credal state is *possible*. It's not clear that it is and it is not possible on some theories of credence.

ments that show Caie's Accuracy Measure is incompatible with his views on Rational Introspection *and* leads to undesirable results when considered in Dutch Book arguments.

How we apply accuracy-dominance arguments to situations is also called into question by an example presented by Greaves (2013).[12]

> *Imps*:
> Emily is taking a walk through the Garden of Epistemic Imps. A child plays on the grass in front of her. In a nearby summerhouse are ten further children, each of whom may or may not come out to play in a minute. They are able to read Emily's mind, and their algorithm for deciding whether to play outdoors is as follows. If she forms degree of belief $x = 0$ that there is now a child before her, they will come out to play. If she forms degree of belief $x = 1$ that there is a child before her, they will roll a fair die, and come out to play if and only if the outcome is an even number. More generally, the summerhouse children will play with chance $(1 - 0.5x)$. Emily's epistemic decision is the choice of credences in the propositions Cr_0 that there is now a child before her, and, for each $j = 1, \ldots, 10$, the proposition Cr_j that the jth summerhouse child will be outdoors in a few minutes' time. (Greaves, 2013, p.918)

In this example, Emily is offered an epistemic bribe, if she can somehow make herself disbelieve something she has clear evidence for

[12] Notation altered for consistency.

then she can maximise the accuracy of her beliefs overall. If she has $Cr_0 = 1$ then the chance of each summerhouse child coming out to play is 0.5 so Emily ought to have $Cr_j = 0.5$, this is the best she can do. Either the child will come out to play (1) or they will not (0) and $Cr_i = 0.5$ is halfway between these. If on the other hand, she has $Cr_0 = 0$ then she can have $Cr_j = 1$ and she has the guarantee that her beliefs match the truth perfectly. If we take accuracy to be the primary epistemic goal of the agent then this suggests that she take the bribe. However, if she takes the bribe then she must ignore/reject an obvious truth. This seems to go against our intuitions about what the agent ought to believe.

Konek and Levinstein (2017) offer a possible solution which rests on a distinction between *epistemic states* and *epistemic acts*. An epistemic act has a causal impact on the world, and we should therefore assess it in terms of its causal impact. An epistemic state is assessed in terms of its fit to the world. They argue that we should be concerned with epistemic states, not epistemic acts. If we adopt this distinction then it has an impact on how we assess self-referential situations. We may be able to see that a certain epistemic act would maximise an agent's accuracy without recommending that the agent change their degree of belief to fit this.

These are some suggestions against the technical framework of Caie's argument. Moreover, we can see that even if we accept that there may be reasons to accept rational probabilistic incoherence in a theory of rationality this recommendation for Hiro does not align with any intuitions we have about the example.

$$(\#) \quad \neg Cr\rho(\#) \geq 0.5$$

When $\neg Cr\rho(\#) \geq 0.5$ it follows that (#) is true and so it follows that Hiro ought to have $Cr\rho(\#) \geq 0.5$. However, when $Cr\rho(\#) \geq 0.5$ it follows that (#) is false and so Hiro ought to have $Cr\rho(\#) < 0.5$ i.e $\neg Cr\rho(\#) \geq 0.5$. If Hiro adopts any precise credence it seems he will have an unstable attitude towards (#).

Although this instability is undesirable one might argue that it is a permissible state for Hiro to be in if he is a real agent. If Hiro is a real agent and is actively reflecting on his credence towards (#) and then updating his credence in light of the changed truth value, then changing his credence in light of new evidence looks desirable. Indeed, continuously updating one's credence in light of changing evidence seems to be exactly what a rational agent ought to do. While this argument may hold some sway for *real* agents it seems that it ought to be possible for an ideal agent to be aware of the set up of (#) and adopt a stable attitude towards it. However, if we consider an ideal agent it seems like they ought to be able to have a stable belief towards (#) rather than adopting a credence that they know in advance to be unstable. Another option, that avoids Caie's suggestion to accept rational probabilistic incoherence, is to question the model of rationality that we are working with. In particular, we could consider another category beyond degrees of belief: suspended judgment. Considering another category looks like a natural move given the parallels that we can observe between probabilistic self-reference and alethic self-reference. In the following section I will outline the parallels between the Probabilistic Liar and the Liar paradox and why suspended judgment looks like an intuitive response to the Probabilistic Liar.

1.2.2 Parallels to the Liar Paradox

The Liar paradox is a sentence that says of itself that it is false.

(λ) (λ) is false.

The Liar paradox is problematic as it leads to contradiction if we accept some intuitive features of truth.

One way of addressing the problems generated by the Liar paradox is to consider three-valued valuation schemes. Maudlin gives a theory of truth where the Liar paradox does not get assigned a truth value of 'true' or 'false'. In Maudlin's theory, there are two categories of sentences that we can give a truth value of 'true' or 'false' to. 1) Sentences whose truth value is determined by the truth value of other sentences and 2) sentences whose truth value is not in any way determined by semantic facts but only by "*the world*" (Maudlin, 2004, p.31). In the case of 1) there are clearly some sentences of the language where the truth value is determined by its components, i.e. a sentence that is the conjunction of two other sentences. In this case, we can determine the truth value of the sentence by the truth values of its conjuncts. Maudlin thinks of this dependency relation as a directed graph; so, for example, with the sentence, $A \wedge B$, the truth value is determined by the truth values of the sentences A and B so there would be arrows from A and B to $A \wedge B$. With this directed graph there will be a boundary, and the sentences on the boundary are those that do not have any arrows pointing to them (only away from them). These are the sentences whose truth value is determined by the world.

Sentences which we can uncontroversially give a truth value (of true or false) to are sentences that can trace their path back to boundary sentences. Maudlin considers it a fundamental principle that “truth and falsity are always ultimately rooted in the state of the world” (Maudlin, 2004, p.49). Sentences that are not boundary sentences or cannot be traced back to boundary sentences (such as the Liar or the Truth Teller) are assigned a third truth value of ‘*ungrounded*’.

Looking to parallels between the Liar paradox and the Probabilistic Liar this suggests a way we might try to resolve the issues that the Probabilistic Liar poses. Within an all-or-nothing model of belief where there are no degrees of confidence, only belief or disbelief, it is also natural to consider a third category of suspended judgment. This third category looks like a good candidate for the doxastic parallel category to ungrounded. Since I am looking at degrees of belief this suggests examining how suspended judgment can be viewed in the credal model of belief. I will outline what is meant by suspended judgement and then examine whether this extra category or distinct attitude can be used to resolve the problems generated by the Probabilistic Liar. An immediate worry in taking this approach is that it will be possible to generate a Revenge problem using suspended judgment. I will address how a Revenge problem can be generated and a solution to it in Chapter 4. In the following sections I will discuss what suspended judgment is and show how it helps to solve the original

Probabilistic Liar.

1.3 Suspended Judgment

Suspended judgment (or agnosticism) is a neutral attitude an agent can have towards a proposition. Some make a distinction between suspended judgment and agnosticism. Monton (1998) for example, views that there is a weak sense in which agnosticism is suspended judgment but that agnosticism in a robust sense is to suspend judgment and also believe that one will not get evidence that would decide the matter. I will talk about agnosticism and suspended judgment in the weak sense and use the terms interchangeably.

Suspended judgement is often (and perhaps most naturally) talked about in the context of a traditional belief framework where an agent has belief or disbelief (rather than degrees of belief). In a traditional belief framework it seems that belief and disbelief alone cannot capture the type of attitude we take towards certain propositions, we may want to withhold belief or disbelief or suspend judgement.[13] Given this, it is useful to first look at what we take suspended judgement to be in the traditional framework. This will inform how to understand the attitude in the context of a degree of belief framework. In this section I will examine how we can understand suspended judgment and in section 1.4 I will look at how we can represent suspended judgment in the context of degrees of belief.

It is not clear that a reductive analysis of suspended judgment can be

[13] See for example Bergmann (2005).

given. There are still many things we can say about when the suspended judgment attitude is adopted and how it differs from credences or belief attitudes. One thought is that suspended judgment can be reduced to non-belief or to non-belief plus some additional criteria. If this is the case, then suspended judgment is not an attitude itself. It seems however that attempts at accounts that reduce suspended judgement to non-belief in some way are inadequate at capturing suspended judgment (see Friedman (2013b) and Sturgeon (2010)). In particular, trying to identify suspended judgement with non-belief (plus some additional criteria) fails to capture its neutrality.

Suspended judgment is an attitude that is arrived at after considering a proposition; it is not simply a catch-all category. This sets it apart from merely not believing. There are lots of things about which we don't have beliefs about, such as propositions that we have never thought of or propositions we cannot understand. These cases seem distinct from the sorts of cases where agents suspended judgement. Having considered a proposition is also important since an agent must at least grasp a proposition in order for them to suspend judgment on it (Friedman, 2013b, p.168). I do not have any attitude towards how good a theory string theory is for example. It is not the case that I am suspending judgment on how good a theory it is, I just do not know enough physics to be able to understand what the theory even says. We can see however that grasping a proposition is also not enough to say that an agent suspends judgement on it. If we are considering the circumstance under which an ideally rational agent should suspend judgment it looks like an agent ought to have considered or deliberated on the proposition. Take the following example where we can see a distinction between two agents that both grasp a proposition but only one has deliberated on it.

> *Alien*:
> Imagine two agent's Alice and Nathan who are both asked about their beliefs concerning the existence of aliens. Alice has never considered this and thus has no belief in favour or against the existence of aliens, it is not the case however that Alice suspends judgment, she just currently does not have a belief about the proposition 'aliens exist'. Nathan on the other hand has thought about whether aliens exist and has decided that there is insufficient evidence to decide either way, thus until further evidence can be gathered he suspends judgment about the proposition 'aliens exist'.

The fact that Nathan has considered and deliberated on the proposition and Alice has not appears to be what gives the intuition here that Nathan is suspending judgment and Alice is not.

In the *Alien* example above we can also see that it is not deliberation alone that gives us the intuition that Nathan is suspending judgment. As well as having considered the proposition he has also finished deliberating on it for the time being and thought that there was insufficient evidence. Friedman points out that we can imagine that an agent has started a deliberation process and got distracted partway through. In such a case even though the agent can grasp the proposition and has started deliberating, it seems they are not yet in a state of suspended judgment (i.e the agent may later finish deliberating and decide that they, in fact, believe aliens exist).

1.3.1 Norms for suspended judgment

What are the circumstances under which it is rational to adopt the suspension of judgment state? It seems clear that there are many circumstances under which it would *not* be rational to adopt it, such as cases where you are certain a proposition is true, or cases where you are certain of the likelihood of a particular event happening (such as your degree of belief in a fair coin landing heads).

From the *Alien* example we can see that there are certain conditions required to recommend or endorse the suspended judgment category. It seems like deliberation or consideration of the proposition is important in distinguishing between an agent who suspends belief and an agent who has no belief. This suggests the norm:

> **Deliberation:** If an agent has not considered / deliberated on a proposition they should not suspend judgment.

Friedman (2017) goes further and argues that inquiry or having an 'interrogative attitude' [14] is both necessary and sufficient for suspending judgment. As interrogative attitudes are typical of inquiry she concludes that the reason we suspend judgement is because of inquiry. Suspended judgment can therefore be an attitude that is arrived at after deliberating, but not a final attitude that is reached in inquiry (since suspended judgment leaves open the possibility of continuing to inquire and settling on belief or disbelief).[15]

[14] Interrogative attitudes are taken to be question-directed attitudes.

[15] Staffel (2019a) argues that being a transitional attitude and being able to change your attitude are key features of suspended judgment. Belief and disbelief in comparison are only appropriate as terminal attitudes.

Suspension of judgment also seems like the appropriate attitude to adopt when there is inconclusive evidence. Feldman, for example, says

> ... if a person is going to adopt any attitude toward a proposition, then that person ought to believe it if his current evidence supports it, disbelieve it if his current evidence is against it, and suspend judgment about it if his evidence is neutral (or close to neutral). (Feldman, 2000, p.679)

This suggests the norm:

> **No Evidential Support:** If the evidence for a proposition does not provide support for belief or disbelief about a proposition then an agent may suspend judgment.

If we take adhering to evidence as one of our epistemic virtues, then in situations where the evidence does not promote a particular belief, suspension of judgment seems like the best option. The **No Evidential Support** norm is a particular instance or a wider norm:

> **No Violation of Norms:** If having an attitude other than suspended judgment leads to another epistemic norm being violated then an agent should suspend judgement.

In the case of evidential support, we regard adhering to the evidence as an epistemic virtue. This can be included in a theory of rationality more formally as the norm:

> **Evidence:** A rational agent should respect the evidence.

If this is included in our theory of rationality as a norm, then if an agent has insufficient evidence for a credence in a proposition they should respect the evidence and suspend judgment. This leaves a wide range of cases where an agent may have insufficient evidence to form a belief, and we can see that there are also a number of ways in which an agent can have insufficient evidence to form a credence. There might be no relevant evidence available (such as with a scientific theory where we are awaiting relevant experimental evidence). There may be evidence, but the evidence is equally balanced to the likelihood of the truth of the proposition or likelihood of the falsity of the proposition. This could be because of incomplete evidence or conflicting evidence. This might occur in cases where an agent has evidence to be confident in two different propositions but also has evidence that the propositions conflict. An example of this is scientific theories; we might have evidence that theory A and theory B both give good predictions and are considered the best working theories for their respective fields. However, we also know that theory A and theory B contradict each other. In this situation we have strong pragmatic reasons to continue using both theories, this does not mean however that an agent has to be confident in the truth of both or either theory. In this situation, it seems like suspending judgment on either or both these theories is a reasonable stance to have.

We have some norms for the scenarios in which an agent is permitted to suspend judgment. We can also ask how the suspended judgment attitude should interact with an agent's other attitudes. In the next section I outline some coherence conditions for the suspended judgment attitude in a belief framework. In section 1.4.1 I will use these to outline some limited coherence conditions for suspended judgment in a degree of belief framework.

1.3.2 Coherence Principles in a Belief Framework

If a rational agent has suspended judgment on a proposition or propositions, then how should their attitude interact with their attitudes about other propositions? For example, if a rational agent has suspended judgment on two *independent atomic propositions* A and B what should we say about $A \wedge B$ or $A \vee B$? The following look like plausible principles as starting points:

> $\neg$
>
> If you suspend on A then you should suspend on $\neg A$.

If an agent has a neutral attitude towards a proposition then it seems there should be symmetry and they should also have a neutral attitude towards the negation of that proposition.

> $\wedge$
>
> If you suspend on A and suspend on B (where A and B are independent propositions) then suspend on $A \wedge B$.

In the case of conjunction, if the agent has reason to suspend judgment on each of the conjuncts, for example, they have insufficient evidence for belief in A and belief in B then it follows that they will also have insufficient evidence to believe in $A \wedge B$.

> $\vee$
>
> If you suspend on A and suspend on B (where A and B are independent propositions) then suspend on $A \vee B$.

Again, if an agent has insufficient evidence for belief in A and belief in B then it follows that they will also have insufficient evidence to believe in $A \vee B$.

The disjunction and conjunction principles above seem plausible when considering disjunctions or conjunctions of a small number of propositions. There is a worry that this might not naturally extend to a larger number of propositions as it may be possible to generate a type of preface paradox for suspended judgment.[16] That is, it may seem natural for an agent to suspend judgment that the next coin flip of a biased coin will land heads, but it is not obvious that you should suspend judgment on ‘the next 100 coin tosses will land heads’.

Note the requirement that the propositions A and B are *independent* is important. If the propositions are not independent then we can consider the following cases: an agent suspends judgment on A, by the negation condition it follows that she also suspends judgment on $\neg A$. What attitude should she have towards $A \vee \neg A$? Even though the agent suspends judgment in both A and $\neg A$ it seems she ought to believe $A \vee \neg A$. Similarly, we can consider the conjunction of $A \wedge \neg A$. In this case, it seems the agent ought to disbelieve the conjunction.

This gives us some idea of the nature of suspended judgment as talked about in all-or-nothing belief models where we talk about an agent having belief or disbelief. We can also ask whether the suspended judgment attitude makes sense in a degree of belief or credal framework and if it does, what is suspended judgment in a credal context.

[16]See Makinson (1965).

1.4 Suspended Judgment and Degrees of Belief?

It is tempting to suggest that suspended judgment can be accommodated by the credal model by some middling credence (credence 0.5). This looks like the most neutral credence an agent could adopt. It is also a symmetric attitude (if I have credence 0.5 in a proposition A then I also have credence 0.5 in its negation) which seems in line with intuitions about suspended judgment from the traditional belief framework. Indeed, it has been suggested by Hájek that we do not need a separate, distinct category of 'suspension of belief' since "we could generally associate agnosticism with 'middling' probability assignments, belief with 'high' probability assignments, and disbelief with 'low' probability assignments" (Hájek, 1998, p.204). When we look more closely at this suggestion however, we can see that having a middling credence of 0.5 does not capture the neutrality or commitment to neutrality that the suspended judgement attitude has.

It seems like an account of degrees of belief should make a distinction between having degree of belief 0.5 and having a neutral attitude, since we can see such a distinction in an agent's betting behaviour. It is reasonable to assume that there is some correspondence between degrees of belief and betting ratios. A strong version of this claim is the betting interpretation of degrees of belief which claims an agent's degrees of beliefs should be understood in terms of the agent's betting behaviour. This gives rise to the well known Dutch Book arguments for Probabilism. If an agent accepts a series of bets that will lead to a sure loss for them then they have been Dutch Booked. If an agent's degrees of beliefs are not probabilistically coherent then they are susceptible to being Dutch Booked. On this interpretation if an agent has degree of belief 0.5 in a proposition *P* this does not mean that they are suspending judgment, rather it means they would be willing to buy or sell a bet for 50*p* if the bet pays £1 if *P* and £0 if not *P*. There are a number of objections to betting interpretations of degrees of belief, so one might think the above example is too simplistic a picture to explain what we take credence 0.5 to represent about an agent's attitudes.

However we do not need to endorse such a strong view of the link between betting behaviour and degrees of belief, merely thinking there is a general and rough correspondence between degrees of belief and betting ratios is enough to show that 0.5 cannot represent a neutral attitude. A further reason that interpreting degree of belief 0.5 as suspension of judgment is problematic can be found by looking at the connection between chance and credence. It seems reasonable to adopt a norm of rationality that gives some connection between chance and credence. A number of chance-credence principles have been given, such as Lewis's Principal Principle. The Principal Principle essentially says that a rational agent's beliefs should conform to the chances. More formally it says:

> ... the following is true. Assume we have a number x, proposition *A*, time t, rational agent whose evidence is entirely about times up to and including *t*, and a proposition *E* that (a) is about times up to and including *t* and (b) entails that the chance of *A* at *t* is *x*. In any such case, the agent's credence in *A* given *E* is *x* ... (Weatherson, 2016a, §5.1).

This suggests that a rational agent should have degree of belief 0.5 in a coin flip landing heads. Having degree of belief 0.5 in this situation is not an indication of neutrality or suspension of judgment about whether the coin will land heads. An account of suspension of judgment should distinguish between situations where an agent suspends judgment about a proposition and situations where an agent has the credence 0.5 in a proposition on the basis of evidence for the chance being 0.5.

In general, if we took the credence of 0.5 as representing the attitude for suspending judgment we would hide the distinction between the balance and weight of evidence. That is, it looks like there ought to be a distinction made between your attitude in cases where a fair coin is flipped and you have a credence in 'the coin will land heads' and cases where a coin of unknown bias is flipped and you have an attitude in 'the coin will land heads'. In the case of the coin with unknown bias, if we are in a precise probability framework, then adopting credence 0.5 is the best we can do to represent that an agent is suspending judgment. There are accounts that explain the distinction by appeal to resiliency (Skyrms (1977)) or stability (Leitgeb (2014)) of belief under updating by conditionalisation. The idea is that if you saw a few flips of a fair coin this is a resilient or stable belief (because of the weight of evidence you have for it) and so you wouldn't change your credence that the 'next coin toss will land heads' (even if the last 10 flips have been heads). In contrast, in the biased coin cases where we want to represent that an agent is suspending judgment, if you do some trials of the coin you are likely to change your credence in 'the coin will land heads'.

This looks like it goes some way to providing the distinction in these cases, however, it still looks like 0.5 is problematic since in cases where the agent is not getting new evidence there is no clear distinction between cases where an agent has credence 0.5 in a fair coin landing heads and cases where an agent has an attitude towards a coin of unknown bias landing heads. Moreover, we can see that given certain assumptions about coherence conditions for suspended judgment in a credal context, it looks like suspended judgment cannot be straightforwardly reduced to a middling credence. In the following section I consider what coherence principles for suspended judgement in a degree of belief framework might look like and why credence 0.5, therefore, looks problematic.

1.4.1 Coherence Principles in a Degree of Belief Framework

Suppose that an agent, Scarlett, suspends judgment towards a number of propositions including towards A_1: a coin$_1$ of unknown bias landing heads, and A_2: a coin$_2$ of unknown bias landing heads, and A_3: a coin$_3$ of unknown bias landing heads. As noted in section 1.3.2 it seems natural that for a small number of independent propositions if Scarlett suspends judgment in A_1, A_2 and A_3 then she would also suspend judgment in the conjunction $(A_1 \wedge A_2 \wedge A_3)$ and disjunction $(A_1 \vee A_2 \vee A_3)$ of these propositions. This intuition provides further support for the claim that suspended judgment is distinct from having a middling credence.

If we take suspended judgement to be represented in a credal framework by a middling credence (a credence of 0.5) then if A_1,A_2 and A_3 are probabilistically independent and we let Cr_S refer to Scarlett's credence we have the following:[17]

$$\begin{aligned} Cr_S(A_1 \wedge A_2 \wedge A_3) &= Cr_S(A_1)Cr_S(A_2)Cr_S(A_3) \\ &= 0.5 \times 0.5 \times 0.5 \\ &= 0.125 \end{aligned}$$

So her credence in $Cr_S(A_1 \wedge A_2 \wedge A_3)$ is not a middling credence.

$$\begin{aligned} Cr_S(A_1 \vee A_2 \vee A_3) &= Cr_S(A_1) + Cr_S(A_2) + Cr_S(A_3) - \\ &\quad Cr_S(A_1 \wedge A_2) - Cr_S(A_1 \wedge A_3) - Cr_S(A_2 \wedge A_3) \\ &\quad + Cr_S(A_1 \wedge A_2 \wedge A_3) \\ &= 0.5 + 0.5 + 0.5 - (0.5 \times 0.5) - (0.5 \times 0.5) - (0.5 \times 0.5) \\ &\quad + (0.5 \times 0.5 \times 0.5) \\ &= 0.875 \end{aligned}$$

Again we see her credence in $Cr_S(A_1 \vee A_2 \vee A_3)$ is not a middling credence. This suggests that suspended judgment cannot be straightforwardly reduced to a middling credence since it seems reasonable that, at least for a small number of propositions, if an agent suspends judgment in several propositions A_1, A_2, A_3 they should also suspend judgment in the conjunction and disjunction of these propositions.

How suspended judgment and credences interact also looks like a difficult question. We can however say something about the extremal cases of credences at least (i.e. credence's of 0 and 1). The following

[17] Adaption of an example from Friedman (2013a).

look like plausible principles:

$\wedge$

If you suspend on A and have credence c in B then:

- If $c = 1$, suspend on $A \wedge B$.
- If $c = 0$, have credence 0 in $A \wedge B$.

$\vee$

If you suspend on A and have credence c in B then:

- If $c = 1$, have credence 1 in $A \vee B$.
- If $c = 0$, suspend on $A \vee B$.

If we only consider extremal cases (or an all-or-nothing belief model) then another natural thing to say about suspended judgment is that an agent cannot simultaneously believe or disbelieve a proposition and suspend judgment about it. If an agent believes a proposition they are confident that it is true, to simultaneously have a neutral attitude towards the proposition seems at odds with this. Hájek (1998) makes a similar point noting how strange it would be if someone were to say "I'm agnostic regarding the existence of quarks; but I believe they don't exist" (1998, p.203).

This suggests that in the credal model we should say something similar. If you suspend judgment about a proposition A you should not simultaneously have a credence c in that proposition for any credence c. As in the all-or-nothing belief model, there is an intuition here that an agent cannot simultaneously have a neutral attitude about a proposition while also having a specific degree of confidence in that proposition. This feature in particular seems important for the aim of using suspended judgment as a solution to the Probabilistic Liar. As I showed in section 1.2 it is problematic if the agent Alex adopts any precise credence in (α). Suspended judgment looks like the natural attitude to consider as an alternative attitude Alex could adopt towards (α). Since a natural intuition is that an agent cannot simultaneously suspend judgment and have a precise credence in a proposition, this makes suspended judgment an ideal attitude to consider as a solution. In the following section I will present an alternative way we can understand suspended judgment in a credal context: as imprecise credences. After presenting an account of suspended judgment in terms of imprecise credences I will return to the above coherence conditions for suspended judgment in a credal context and show how my account meets these constraints.

1.5 Suspension as Imprecise Credence

How can an agent suspend judgement or have a neutral attitude towards a proposition *and* have a credence in that proposition? I have already noted that there are problems with thinking of suspended judgment with credence 0.5. However, if we assume that in a credal model an agent can only have precise credences, 0.5 does seem like the best we can do in an attempt to express a neutral attitude.

This suggests looking outside the traditional sharp credal framework. One solution is to treat the suspended judgment attitude as expressing that the agent has indeterminate or imprecise credences. There is a range of terminology in the literature that refers to a number of related concepts including referring to indeterminate credence, interval-valued credence or mushy credences. I will follow S. Bradley (2016) and use the term imprecise as an umbrella term. Imprecise probabilities represent an agent's belief by a set of credence functions rather than a single probability. Rather than having some precise credence $Cr(x)$ in proposition x we represent the agent's belief by a credal set[18] or representor[19] C which is a set of probability functions.

$$C = \{Cr_1, Cr_2, Cr_3, \dots\}$$

Where the Cr_i are credence functions. We assume each of the credence functions in the credal set is probabilistically coherent.

Joyce (2010) gives one way we can think about the credal set in terms of a credal committee, each credence function in the credal set represents the credence of a member of the committee. Together the credence functions in the set represent the opinion of the committee. In the case where C contains only one credence function, we are in the trivial (precise) case. We can also give a summary function for the credal set for an agent's belief in a proposition x:

$$C(x) = \{Cr(x) : Cr \in C\}$$

We can also talk about the upper and lower probabilities of an agent's degree of belief in x. We can write $C(x) = [a, b]$ where $a, b \in [0, 1]$ and

[18] In the terminology of Levi (1974).

[19] In the terminology of van Fraassen (1990).

$a = \inf C(x)$ and $b = \sup C(x)$. This gives us some information about the range of the credences in the credal set but it is not a fully adequate representation. If we can have two agents A and B that both have credal sets that range over $[0.2, 0.6]$ but agent A has an even spread of credences across this range and agent B has a concentration of credences at the upper end of this range then the summary function will not let us know this difference. However, for the purposes of identifying general motivations for adopting imprecise credences and identifying features of imprecise credences that seem to make it appropriate for the suspended judgment attitude it will suffice to refer to the upper and lower probabilities.

There are a number of different accounts/ interpretations of imprecise probabilities and different justifications or reasons for taking these different interpretations.[20] In Chapters 2 and 3 I will defend in detail a particular interpretation of an imprecise credal set for the suspended judgment attitude. For now, in this chapter, I will outline the general reasons why we might want to adopt imprecise credences and why imprecise credences seem like a natural way of representing suspended judgment in a degree of belief framework. I will then consider how suspended judgment might be represented within an imprecise credence model.

[20]See Walley (1991), Hájek (2003), Nehring (2009), Kaplan (2010), Alon and Lehrer (2014) and Rinard (2017).

1.5.1 Why Imprecise Credences ?

Imprecise probabilities (or imprecise credences) have been argued for or posited for a number of reasons. In this section I will outline examples that suggest why, regardless of suspended judgement, we ought to adopt a imprecise credal framework. In section 1.5.2 I will then argue for why suspended judgment ought to be represented by imprecise credences in a credal framework.

There are several reasons why precise probabilities seem like too restrictive a model to represent agents' attitudes. One motivation for adopting imprecise credences is due to scenarios where an (ideal) agent does not have sufficient evidence to justify forming a precise credence. For example, in the Ellsberg problem which is a situation where the agent knows the following:

> There is an urn that contains 90 balls, 30 of the balls are red and the remaining balls are a mix between blue and yellow (the agent has no further information about the proportion of blue and yellow).

If the agent is asked to bet on whether a ball selected from the urn is yellow they have insufficient evidence to have a precise credence. With imprecise credences we can describe the agent's credences as $Cr(R) = 1/3$ and $C(B) = C(Y) = [0, 2/3]$. An imprecise credence model can help make sense of the way real agents act in decision situations as well. In the Ellsberg example, agents are put in a decision situation where they are offered bets about risky propositions and bets about ambiguous propositions. A risky proposition is one where there is a known chance whereas it is ambiguous if the chances are not known or only partially known as in the situation I outlined. We see that when real agents are offered bets they tend to act in a certain way: ambiguity aversion.[21] This betting behaviour can be better explained if we attribute imprecise credences to the agent and the view that the agent has vague or indeterminate attitude towards the ambiguous proposition.[22]

Joyce also motivates imprecise credences in terms of lack of evidence saying:

> Imprecise credences have a clear epistemological motivation: they are the proper response to unspecified evidence (Joyce, 2005a, p.171)

The connection between evidence and precision of credence is also emphasised by Sturgeon (2008) who says that if there is sharp evidence then a rational agent should have a precise credence but in circumstances where the evidence is "essentially fuzzy ... it warrants ... a fuzzy attitude" (p.159).

One might think in these cases the agent has vague evidence for a proposition, but that, given further evidence the agent could in principle come to have a precise credence. Another reason one might adopt imprecise credences is if the proposition the agent has an attitude towards is itself vague. Then it is natural to reject the idea that

[21] See Camerer and Weber (1992).

[22] See Mahtani (2019a) for an overview of different decision rules that have been suggested for imprecise credences.

an agent should have a precise credence towards it. It seems like an agent ought to have a vague or indeterminate attitude towards vague propositions, if we take cases where 'Harry is tall' is borderline it seems like an agent ought not to have a precise credence towards 'Harry is bald'.[23] Lack of evidence or vague evidence as a motivator for imprecise credences suggests that it is not just the case that it is permissible for an agent to adopt imprecise credences but rather that an ideal agent is rationally required to adopt imprecise credences.

Imprecise credences are also necessary if we want to be able to model or accommodate that an agent has incomplete attitudes (I will discuss this motivation in much greater detail in the next chapter). The traditional picture of precise credences requires that agents have complete beliefs, that is to say for any two propositions A, B it is either the case that the agent thinks A is at least as likely as B, B is at least as likely as A, or they are equally likely. If agents have precise attitudes towards all propositions then this will of course follow, however, this seems like an unrealistically strong constraint. In some cases, it seems reasonable that an agent might just be indifferent between the relative likelihood of two propositions. That is, an agent may simply lack an attitude.

1.5.2 Imprecise credences and suspended judgment

As we can see, there is an overlap between the motivations one might have for adopting imprecise credences and the reasons why one might

[23] See Rinard (2017) for an account of imprecise credences motivated by vagueness.

adopt suspended judgment. In both cases, a lack of evidence looks like a key reason why one would adopt imprecise credences (rather than a precise credence) or suspended judgment (rather than belief or disbelief). Walley (1991) identifies several sources of imprecision including lack of information, conflicting information, conflicting beliefs, information of limited relevance and physical indeterminacy (pp.212-215) all of which are compatible with intuitive ascriptions of suspended judgment. Friedman (2013a) also considers imprecise credences as an interpretation of suspending judgment noting it is a possible solution for someone "who wanted to avoid the conclusion that suspension is rationally compatible with any precise degree of belief"(p.75).

Equating suspended judgment to having imprecise credences in a credal framework, therefore, seems like a natural thought. We can see that by viewing suspended judgment as having imprecise credences we have a way of recommending what attitude an agent should have in various situations where there is a lack of evidence or in cases of uncertainty. Imprecise probabilities can provide a way of distinguishing between situations where the agent does not have sufficient evidence to form a precise credence (and thus wants to suspend judgment) and situations where the agent has evidence to form the precise credence of 0.5. For example, we can consider two scenarios where an agent is being asked their credence in a coin flip landing heads. In the first scenario, the agent has evidence that the coin is fair and so believes $Cr(H) = 0.5$. In the second scenario, the agent has no evidence that the coin is fair and so can instead have the credence $C(H) = [0, 1]$.

That being said, having imprecise credences encompasses a wide range of attitudes. One might, for example, have imprecise credence towards a proposition in the range $[0.8, 0.9]$. This however does not look appropriate as a suspended judgment attitude. In this case, we can say that the agent is confident towards the proposition even if they have imprecise credences. Similarly, if an agent has the imprecise credal set of $[0.1, 0.2]$ towards a proposition is seems they are unconfident. Rather than identifying suspended judgment with having any imprecise credence we should instead identify certain types of imprecise attitudes an agent can have. There are a number of existing accounts that talk about suspended judgment or agnosticism as imprecise credences (see van Fraassen (1998), Hájek (1998), Monton (1998), Sturgeon (2010) and Friedman (2013a)). There is however disagreement about what form of imprecise credences represent suspended judgment.

I will consider an account of suspended judgement in terms of imprecise credences given by van Fraassen (1998). There are a number of features of his view that look desirable in capturing the attitude of suspended judgement. His view however focuses on an attitude he calls 'negative suspension' which I will argue fails to capture the suspended judgment attitude. By adapting some of van Fraassen's ideas I will argue for a general form that imprecise credences should take in order to denote suspended judgment.

van Fraassen argues that we can represent what it means to have no opinion at all in terms of vague subjective probability and this can be represented in terms of the upper and lower probabilities of the interval of the representor of the agent. Having an attitude of 'negative suspension' is represented by the interval $[0, n]$ for some number $n \in [0, 1]$ (where n can be less than 1). For van Fraassen's account, we can see that suspended judgment (or agnosticism as he refers to it) is not a unique attitude. There are a variety of ways that one can have suspended judgement, which might be represented by a variety of intervals (which can be represented by the constraints on the interval being minimal enough that a variety of credal intervals can meet the conditions). He does however have a strict condition on the form of the credal set, we can see that no amount of conditionalising on new evidence can change the form of $[0, n]$. If one of the credence functions in the credal set assigns 0 to a proposition an agent is stuck with it.

van Fraassen takes this as an advantage, he takes suspended judgment (or agnosticism) to be a very robust notion. Once an agent has suspended judgment in a proposition he thinks they cannot come to believe it (or be confident in it). The form he gives does however leave open the possibility that an agent could come to disbelieve something they used to suspend judgment in. This analysis is very different from my intuitions about suspended judgment.[24] It seems that, if it is possible for an agent to come to disbelieve something they originally suspended judgment in then it ought to also be the case

[24] One explanation for the distinction in intuitions is that we may be using the term suspended judgment or agnosticism differently. van Fraassen is discussing agnosticism in the context of views about scientific realism. He does however also relate his use of the term agnosticism to its use in the context of belief about God's existence suggesting that his use of the term agnosticism is supposed to latch onto a more general category than just scientific agnosticism.

that the agent could come to believe something they originally suspended judgment in. For an imprecise credal set to represent suspended judgement there should not be a requirement that 0 is in the credal set.

As well as a difference in intuitions we can also see that if it is required that 0 is in the imprecise credal interval it seems problematic when we consider the question of whether one thinks that suspended judgment ought to be a symmetric attitude in a degree of belief context. In a traditional framework, it seems natural to say that if one suspends judgment A then one ought to also suspend in not-A.[25] In a probability framework we can say if an agent has some precise or imprecise probability to A then she ought to have the same precise or imprecise probability to $\neg A$. In a precise credal framework the only credence that one could assign is 0.5. As I have already argued however, this seems like it fails to capture the suspended judgment attitude. In an imprecise credal framework, if we take on van Fraassen's requirement that suspended judgment is an imprecise credal set that includes 0 then the only symmetric attitude is the full interval $[0, 1]$.

The full interval seems like it can capture some instances of suspended judgment. However, a worry with identifying suspended judgment as $[0, 1]$ is that an agent can never change from this attitude. If an agent receives new evidence and updates on this evidence (by conditionalising) then their interval will remain $[0, 1]$ even if the spread of credences within the interval changes. I am making the assump-

[25] Hájek (1998) also makes this point and argues that suspended judgment is a symmetric attitude.

tion that when an agent with imprecise credences updates her beliefs she updates each credence function in the credal set, and therefore the credence functions that start at 0 and 1 are stuck there.[26] S. Bradley (2016) refers to this as the problem of belief inertia.[27] As noted above it seems like an agent ought to be able to sometimes learn and change their attitude from suspended judgment.

There are a number of responses to the belief inertia problem, one is to accept that if one has such little evidence that they are in the maximally imprecise credal set then it is appropriate that they cannot update from this. Walley gives an argument to this effect saying that maximum imprecision is a vacuous state and that

> If the vacuous previsions are used to model prior beliefs about a statistical parameter, for instance, they give rise to vacuous posterior previsions… However, prior previsions that are close to vacuous and make nearly minimal claims about prior beliefs can lead to reasonable posterior previsions. (Walley, 1991, p.93)

Joyce makes a similar point.[28] This seems like a possible response if we think having maximally imprecise credences are a response to only a limited type of proposition, however, it does not seem useful

[26] Joyce gives the update rule that:

> If a person in a credal state C learns that some event D obtains (and nothing else), then her post-learning state will be $C_D = \{Cr(\cdot|D) : Cr \in C\}$ (Joyce, 2010, p.287).

[27] Also discussed not under this name in (Levi, 1980, §13.2) and Walley (1991).

[28] See (Joyce, 2010, p.291).

if one wants to identify suspended judgment with maximum imprecise credences.[29] It seems one should be able to change their mind about something they currently suspend judgment in.

Recall the coherence conditions I presented for suspended judgement in an all-or-nothing belief framework (in section 1.3.2). If an agent suspends judgment in two independent propositions A and B it looks like she also ought to suspend judgment in $A \wedge B$ and $A \vee B$. Moreover, these principles look realistic for conjunctions and disjunctions of small numbers of propositions but not necessarily if we were to consider conjunctions and disjunctions of a large number of propositions. As I showed in section 1.4.1, credence 0.5 does not look promising as a representation of suspended judgment since it fails to meet the conjunction and disjunctions principles even for small numbers of propositions. On the other hand, if we consider the full imprecise credence interval $[0, 1]$ we can see that if one has $C(A) = [0, 1]$ and $C(B) = [0, 1]$ we can see that if we take the minimum of each of these

[29] For an overview of Joyce's (2010) response to belief inertia and problems with it see Vallinder (2018). Vallinder argues that even when agents do not have maximally imprecise credence they can still have belief inertia and that belief inertia is, therefore, a problem for all imprecise credal accounts. Moss (forthcoming) argues that Vallinder is making an implicit assumption about how rationality constraints apply to imprecise credences.

credal sets, $Cr_{min}(A) = 0$ and $Cr_{min}(B) = 0$):

$$\begin{aligned} Cr_{min}(A \wedge B) &= Cr_{min}(A)Cr_{min}(B) \\ &= 0 \times 0 \\ &= 0 \end{aligned}$$

and the maximum of each of these credal sets, $Cr_{max}(A) = 1$ and $Cr_{max}(B) = 1$:

$$\begin{aligned} Cr_{max}(A \wedge B) &= Cr_{max}(A)Cr_{max}(B) \\ &= 1 \times 1 \\ &= 1 \end{aligned}$$

so the full imprecise credence interval would also represent the range of credences one has towards $(A \wedge B)$, i.e $C(A \wedge B) = [0, 1]$.

Similarly we get:

$$\begin{aligned} Cr_{min}(A \vee B) &= Cr_{min}(A) + Cr_{min}(B) - Cr_{min}(A \wedge B) \\ &= 0 + 0 - 0 \\ &= 0 \\ Cr_{max}(A \vee B) &= Cr_{max}(A) + Cr_{max}(B) - Cr_{max}(A \wedge B) \\ &= 1 + 1 - 1 \\ &= 1 \end{aligned}$$

so the full imprecise credence interval would also represent the range of credences one has towards $(A \vee B)$, i.e $C(A \vee B) = [0, 1]$. If $[0, 1]$ represents that an agent suspends judgment then the agent will remain suspending judgment in the conjunctions of propositions she suspends judgment about and in the disjunctions of propositions she suspends judgment about. In particular, we can see that even if we consider conjunctions or disjunctions of a large number of propositions then an agent's attitude would still remain $[0, 1]$. As noted above, this seems unrealistic. Thus lending further support to the full interval being too strict a requirement to adequately represent suspended judgement.

We can avoid the problem by removing van Fraassen's requirement that the credal set includes 0 or even any value close to 0. That is, we leave open the possibility that the lower probability of the interval could be 0.1 or 0.2 or 0.3 etc. Removing the requirement of 0 being in the credal set does however leave us with no constraints on what types of imprecise credences represent suspended judgment. By looking at the above discussion about van Fraassen's account there are certain conditions we can suggest.

We can instead identify suspended judgment with an interval $[a, b]$ $(a \neq b)$ where a does not need to be 0 and b does not need to be 1. While including 0 in the credal set does not seem necessary to represent suspended judgment it does seem necessary that $a < 0.5$ and $b > 0.5$. If all of the credence functions in the set assign a credence of greater than (or less than) 0.5 then it looks like the credal set expresses that the agent is confident (or unconfident) in the proposition, even if this confidence (or unconfidence) is spread out. As noted above, it seems like suspended judgment ought to be a symmetric attitude, if an agent suspends judgment in A then they also ought to suspend judgment in $\neg A$. There are two ways we could characterise a symmetry condition on intervals of the form $[a, b]$.

One way we could require symmetry is to say if an agent has $C(A) = [a, b]$ *and* she suspends judgment in A then $C(\neg A) = [a, b]$. This will be the case whenever $a + b = 1$. Then credal sets such as $[0.2, 0.8]$ or $[0.1, 0.9]$ meet the above criteria. This symmetry condition seems too restrictive, however. If an agent suspends judgment in A and has $C(A) = [0.2, 0.8]$ and suspends judgment on B and has $C(B) = [0.1, 0.9]$ then she will not suspend judgment on $A \wedge B$ or $A \vee B$. That is, her attitude in the conjunction or disjunction will fail to be symmetric in the way outlined above.

Indeed, such a restrictive symmetry condition only seems to make sense if we assume suspended judgment has a unique representation (such as $[0, 1]$). Once we remove the assumption that suspended judgment is represented by a unique credal interval there is another way an agent's attitudes in A and $\neg A$ can be symmetric. Let suspended judgment be understood as an interval $[a, b]$ ($a \neq b$) where a does not need to be 0 and b does not need to be 1 and $a < 0.5$ and $b > 0.5$. Then the following two intervals $[0.2, 0.7]$ and $[0.3, 0.8]$ meet these conditions. Moreover, note that if an agent suspended judgment in A and has the credal set $C(A) = [0.2, 0.7]$ then her attitude towards $\neg A$ is $C(\neg A) = 1 - C(A) = [0.3, 0.8]$. So, the agent suspended judgment in both A and $\neg A$. On this definition of suspended judgment it is also possible for an agent to suspend judgment in two independent propositions A and B and suspend judgment on $A \wedge B$ or $A \vee B$ also. However, unlike with the interval $[0, 1]$ it does not follow that an agent will be stuck suspending judgement when they consider the conjunction of a large number of propositions (where they suspend judgment in each of the conjuncts).[30]

On my characterisation, the interval $[0, 1]$ is also a representation of suspended judgment, but it is not the only credal interval that represents suspended judgment. In rejecting van Fraassen's requirement that suspended judgment is an imprecise credal set that contains 0 we leave open the possibility that in some cases suspended judgment could be a robust state that the agent cannot move from (in the maximum imprecision case). In other cases, the agent can update on her credences.

Like van Fraassen (1998) I take it that the suspended judgment attitude can be represented by a variety of intervals rather than some particular interval. An agent can therefore learn and update on new evidence and come to either be confident or unconfident in a proposition they used to suspend judgment about. Unlike van Fraassen I do not think that being in a state of suspended judgment is one that requires having some part of your credal set express no confidence in the proposition. In addition (as I will argue in detail in Chapters 2 and 3) I take a particular interpretation of the imprecise credal set and therefore a particular interpretation of what intervals of the form

[30] Similarly, if they consider the disjunction of a large number of propositions (where they suspend judgment in each of the disjuncts).

$[a, b]$, $a < 0.5$ and $b > 0.5$ represent.

As mentioned above, one motivation for adopting imprecise credences is that an agent may not have complete comparative orderings of their beliefs. We can represent this with an imprecise credal set, and therefore represent that agents can have credence gaps. Credence gaps can occur when we take the imprecise credal set as representing an agent's determinate attitudes. We can represent that an agent lacks an opinion between propositions A and B by having a credal set where at least one credence function in the credal set represents that the agent is strictly more confident in a proposition A than a proposition B and at least one credence function in the credal set represents that the agent is strictly more confident in B than in A. I suggest we should understand the attitude of suspending judgment as having a credence gap with respect to a proposition A and its negation. When an agent has an imprecise credence towards A of the form $[a, b]$, $a < 0.5$ and $b > 0.5$ she will have a credence gap.[31] We can see this with the example of an agent whose credence span the full range $[0, 1]$.

Scarlett has $C_s(A) = [0, 1]$ and therefore $C_s(\neg A) = [0, 1]$. In particular we can see that if a credence function Cr_1 in C_s assigns credence 0 to A then, by Probabilism $Cr_1(\neg A) = 1$, similarly, if $Cr_2 \in C_s$ and

[31] It also looks like the gap between a and b should be 'big'. For example, the range [0.49, 0.51] does not seem to capture the suspended judgment attitude. What counts as a 'big gap' is vague. Since suspended judgment is an attitude normally understood in a much more course-grained belief framework it follows that there may not be a precise characterisation of it in a degree of belief context. One idea may be to point to something similar to the idea that the credal threshold for belief might be context-sensitive. So it might be that a 'big gap' between a and b is also context-sensitive.

$Cr_2(A) = 1$ then $Cr_2(\neg A) = 0$. So, she has a credence gap towards A and its negation. When an agent has a credence gap there is not enough consensus amongst the credence functions in the credal set. This is reflective of the type of situations that would lead to an agent having credence gaps such as lack of evidence for the agent to form an opinion about the relative likelihood of propositions or conflicting evidence. These look like exactly the kind of reasons why an agent would want to suspend judgment. Moreover, with this understanding of suspended judgment, we can see that it is not compatible for an agent to suspend judgment and have a particular sharp credence in a proposition at the same time. This is an important feature of the account when considering suspended judgment as a solution to the Probabilistic Liar and one of the coherence conditions I outlined for suspended judgment in a credal context in section 1.4.1.

Furthermore, we can see that this account of suspended judgment meets the other coherence conditions. As I've already noted above with my characterisation of suspended judgment it is symmetric (in the sense that if one suspends judgment in A then one also suspends judgment in $\neg A$). When suspended judgment is understood as this it also seems to meet the intuitions I outlined in section 1.3.2 regarding coherence conditions for conjunction and disjunction. We can also see that suspended judgment interacts well with extremal credences (as I outlined in section 1.4.1). Consider two propositions A and B where you suspend judgment in A ($C(A) = [a, b]$) and have either credence 0 or 1 in B.

When $C(B) = 1$

$$\begin{aligned} C(A \wedge B) &= C(A)C(B) \\ &= [a, b] \times 1 \\ &= [a, b] \end{aligned}$$

$$\begin{aligned} C(A \vee B) &= C(A) + C(B) - C(A \wedge B) \\ &= [a, b] + 1 - [a, b] \\ &= 1 \end{aligned}$$

When $C(B) = 0$

$$\begin{aligned} C(A \wedge B) &= C(A)C(B) \\ &= [a, b] \times 0 \\ &= 0 \end{aligned}$$

$$\begin{aligned} C(A \vee B) &= C(A) + C(B) - C(A \wedge B) \\ &= [a, b] + 0 - 0 \\ &= [a, b] \end{aligned}$$

Also, as I noted in section 1.3 there is a distinction between *unawareness* and *suspended judgment*. We can see this distinction in our formal model—one is unaware of the propositions which are not in the algebra of her credence function, and suspends judgement on those propositions which are in the domain of that function but are not assigned precise values.

1.6 Suspending Judgment in the Probabilistic Liar

Revisiting the Probabilistic Liar, we have seen that if Alex has any particular precise credence in the range $[0, 0.5)$ then (α) is true and if Alex has any particular precise credence in the range $[0.5, 1]$ then (α) is false. This leads to any credence Alex adopting being problematic. Considering suspended judgment as a solution is also problematic unless the agent can suspend judgment and have an imprecise credence at the same time. I will now examine how considering suspension of judgment as having imprecise credences works in the Probabilistic Liar example.

Let us consider Caie's argument that shows there is a conflict between Rational Introspection and Probabilism again. Caie's argument as applied to my version of the Probabilistic Liar goes as follows:

$$(\alpha) \quad Cr\rho(\alpha) < 0.5$$

As in Caie's argument we can say:

$$(6) \;\; \rho(\alpha) = \rho\text{`}Cr\rho(\alpha) < 0.5\text{'}$$

Versions of the introspection principles Caie uses are:

$$(7) \;\; [Cr\rho(\alpha) < 0.5] \rightarrow [Cr(\rho\,\text{`}Cr\rho(\alpha) < 0.5\text{'}) > 0.5]$$

$$(8) \;\; [\neg Cr\rho(\alpha) < 0.5] \rightarrow [Cr(\rho\,\text{`}\neg Cr\rho(\alpha) < 0.5\text{'}) > 0.5]$$

Which express that an agent is confident about what credences she currently adopts.

Assume:

$$(9)\quad Cr\rho(\alpha) < 0.5$$

Then by (6) $Cr(\rho`Cr\rho(\alpha) < 0.5') < 0.5$
but given our assumption of (7) we get a contradiction.

So it follows that

$$(10)\quad \neg Cr\rho(\alpha) < 0.5$$

Then by (6) we get $\neg Cr(\rho\ `Cr\rho(\alpha) < 0.5') < 0.5$

and by (8) we get that $Cr(\rho\ `\neg Cr\rho(\alpha) < 0.5') > 0.5$.

i.e. of the form $\neg Cr(A) < 0.5$ and $Cr(\neg A) > 0.5$. By Probabilism and $Cr(\neg A) > 0.5$ it follows that $Cr(A) < 0.5$. Which is a contradiction with $\neg Cr(A) < 0.5$.

This argument assumes that Alex has precise credences, however. Caie has made the background assumption that an agent's attitudes can be determinately represented by a single real-valued credence function.We can see Caie makes this assumption since he assumes both Probabilism and Rational Introspection in his argument and then shows that there is a conflict between these principles of rationality. Probabilism is defined in terms of precise probability functions so Caie must have the background assumption of precise credences. [32]

[32] Moreover, when Caie argues for his own solution he gives it in terms of precise credences(see (Caie, 2013, pp.547-548)) and when he is defining his accuracy measure his credence function gives an assignment of real numbers which again

As Caie shows, if we make this assumption then the Probabilistic Liar is problematic if we also accept Rational Introspection and Probabilism. If we consider an agent who suspends judgment on the Probabilistic Liar Caie's argument has no application, since, one of the background assumptions he makes is false. Caie's argument, so construed, is a reductio of the background assumption that agents are determinately representable by real-valued credence functions.

Once we reject Caie's assumption that an agent has determinate attitudes that can be represented with a real-valued credence function the question remains, how can we understand what it means to suspend judgement in (α). When an agent suspends judgment in (α) they have imprecise credences. One option is to consider how the argument applies to each credence function in the credal set. In particular, when Alex suspends judgment she has a credal set $C(\alpha) = [a, b]$, $a \neq b, a < 0.5$ and $b > 0.5$. For simplicity, I will consider the case where Alex has maximally imprecise credences $C\rho(\alpha) = [0, 1]$. This seems like a reasonable starting point for an agent that has no evidence either way and where any level of confidence or lack of confidence in (α) is problematic.

Some of the credence functions in Alex's credal set say that (α) is more likely than $(\neg\alpha)$ and some of the credence functions say that (α) is less likely than $(\neg\alpha)$. We can see that Alex has a credence gap for the pair of propositions (α) and $(\neg\alpha)$. When Alex is in suspense between (α) and $(\neg\alpha)$ this blocks the instability between her thinking first (α) is true and

shows us he is working with a precise credence function (since in contrast imprecise credences give an assignment of sets of reals (see (Caie, 2013, p.544)).

then $(\neg\alpha)$ is true. The uncertainty the agent has about their credence in (α) is carried over to uncertainty about the truth value of (α). In fact, when an agent suspends judgment on (α) it seems that (α) is not true or false but rather falls into some third category such as Maudlin's ungrounded. When an agent suspends judgment there is nothing about their attitude that suggests they ought to adopt $Cr(\alpha) = 1$ or $Cr(\alpha) = 0$. So suspended judgment, and in particular, $C\rho(\alpha) = [0, 1]$ blocks Caie's argument by rejecting the strict definition of Probabilism which requires an agent's attitudes are represented by a single real-valued credence function and replacing it with a weaker version. When the agent introspects on their credences about (α), even if we grant that their credal set is completely transparent to them we do not get the instability that occurs when they have a precise determinate attitude towards (α). This shows how suspended judgment provides a stable attitude towards (α).

Adopting imprecise credences also looks like it can provide a solution to the *Leap* example I gave in section 1.1. Recall in the Leap example, our agent James is in a scenario where he has to leap across a chasm. Whatever James's degree of confidence $x \in [0, 1]$ is, his chance of making the leap is x. In this case, there are conflicting intuitions about what attitude it is rational for James to have. The extremal credences appear to be the most accurate credences James can have. It also seems possible that any degree of belief is permissible since any degree of belief he has lines up with the chances. Given this, it seems one reasonable response James could have in this scenario is to suspend judgment on whether he will succeed in making the jump.

Given the suspended judgment attitude it follows that we reject Caie's assumption that an agent has determinate attitudes that can be represented with a real-valued credence function. Caie's argument as given above does not apply, and we would have to rewrite it in terms of imprecise credences in order to talk about an agent suspending judgment. If we run the argument in terms of precise credences by making the presupposition that the agent has a precise attitude towards the Probabilistic Liar then Caie's argument becomes a reductio of the supposition that the attitude towards the Probabilistic Liar is a precise attitude.

Given this a further question remains:

Can we express a version of the Probabilistic Liar in terms of credal sets?

That is, rather than

$$(\alpha) \quad Cr\rho(\alpha) < 0.5$$

can we consider:

$$(\alpha^*) \quad C\rho(\alpha^*) < 0.5$$

Where C is the credal set Alex has and $C\rho(\alpha^*) < 0.5$ can be understood as expressing that for every credence function $Cr \in C, Cr\rho(\alpha^*) < 0.5$.

(α^*) is a version of a Revenge problem for the Probabilistic Liar. That is, once we have given a solution to the original problem, can we generate a new problem in terms of that solution. I address this question in Chapter 4 where I consider problems that can be generated when a self-referential proposition can refer to the content of an agent's credal set. As well as re-writing the Probabilistic Liar in terms of imprecise credences we would also have to give versions of his Rational Introspection norms and Probablisim in terms of imprecise credences. The Probabilism norm does not hold in the case of imprecise credences, there is however a natural way to give a weakening of Probablisim for imprecise credences. We claim that every precise credence function in the credal set ought to obey Probablisim.[33]

We also have to give versions of Caie's rational introspection principles in terms of imprecise credences. While his rational introspection principles look plausible for precise credences it is unclear whether they are plausible for imprecise credences. In the precise case it seems reasonable that an agent be confident regarding their precise attitude. As I will discuss in more detail in Chapter 4 where I address the Revenge problem, I argue that in the case of Revenge it should be indeterminate what imprecise attitude the agent has towards (α^*). In this case the imprecise versions of Caie's introspection principles seem far too strong. When it is indeterminate whether an agent's attitude towards (α^*) is more or less than 0.5 it seems like too strong a requirement that an agent be confident in their attitude when they introspect (even if it is in fact the case that their attitude is less than 0.5 or more than 0.5).

One immediate worry is that it seems very easy to generate a Revenge problem from my solution to the Probabilistic Liar. In Chapter 4 I discuss the fact that it is so easy to generate a Revenge Problem is a

[33]There may be room here then to avoid the contradiction in Caie's argument if one of the credence functions in a credal set assigns a credence of greater than 0.5 to a proposition A and another credence function in the credal set assigns a credence of greater than 0.5 to $\neg A$.

familiar problem from theories of truth. Once one has given a solution to the Liar paradox it is easy to generate a Revenge problem for the Liar. Since this problem seems pervasive in theories of truth and there are clear analogies between the Liar and the Probabilistic Liar, the fact that it is easy to generate a Revenge problem for the Probabilistic Liar is not in itself a worry for my account.

1.7 Conclusion

Self-reference can occur in a probabilistic setting when we consider situations where an agent is not a pure observer. These situations are problematic for a rational agent as there seems to be no stable belief the agent can have. I have suggested the agent suspend judgment in such a situation. The suspended judgment attitude provides a stable attitude for the agent to have and fits with our intuitions about what type of attitude an agent would have when presented with the Probabilistic Liar. Moreover, given the parallels between alethic self-reference and probabilistic self-reference, considering suspended judgment as a possible solution seemed like a natural move. I have outlined how we should understand suspended judgment in a traditional belief framework and its place as an independent attitude an agent can adopt. Using this I showed how we can represent and understand this attitude in a degree of belief framework.

A key feature of suspended judgment that I put forward was the claim that an agent cannot simultaneously suspend judgment and have any particular sharp credence. A solution to this is to think of suspended judgment as having an imprecise credence. We can also see that treating suspended judgment as having imprecise credence is natural when we consider the conditions for adopting imprecise credences. When suspended judgment is understood as having imprecise credences this provides a solution to the conflict with Probablisim and Rational Introspection that Caie presents. It rejects the implicit assumption he makes that the only determinate attitudes an agent can have are representable with precise credence functions.

In the next two chapters I give a more detailed account and defence of what interpretation of an imprecise credal set we ought to adopt and how to understand suspended judgment in that interpretation. Given the parallels between the Liar paradox and the Probabilistic Liar, it is also natural to question whether there will be a Revenge problem for the Probabilistic Liar. In Chapter 4 I will outline how Revenge problems can be generated and offer a solution also in terms of imprecise credences.

Chapter 2
Imprecise Credences and Comparativism

2.1 Introduction

I have argued that the attitude an agent ought to have towards the Probabilistic Liar is to suspend judgment, where suspended judgment is understood as having imprecise credence. Suspended judgment is also a determinate attitude. My position on suspended judgment therefore requires that imprecise credences are interpreted as representing an agent's determinate attitudes. Traditional Bayesian frameworks use precise, classical probabilities that obey the Kolmogorov axioms. Using precise credences to represent an agent's attitudes can be considered too unrealistically precise, however. Even if we restrict our attention to ideally rational agents there are situations where an agent lacks sufficient evidence to justify the level of precision in their credences the traditional Bayesian account requires.

Imprecise credence models are an alternative framework that has been developed. Imprecise credences represents an agent's attitudes in terms of sets of credence functions rather than one precise credence function. When we represent an agent's beliefs with an imprecise credence model we need an interpretation of the set of credence functions in order to explain which parts of the model represent features of an agent's belief. In turning to an imprecise credal set we are rejecting the claim that precise credence functions represent an agent's beliefs. There are a number of interpretations of imprecise credence models that take different aspects of the mathematical structure as having representational importance.

These different interpretations diverge significantly in what explanatory or predictive results they give. In this chapter and in chapter 3 I will argue for a comparativist intersectionist interpretation. This interpretation supports my claim that the imprecise credal set that represents suspended judgment represents a determinate attitude. In section 2.2 I will outline what this interpretation of imprecise credences is. A comparativist intersectionist interpretation is formed of two parts. The comparativist component rests on a view about comparative beliefs - that we can describe an agent's degrees of beliefs in terms of their comparative beliefs. An intersectionist interpretation of an imprecise credal set is an interpretation that takes the credal set as a whole as representational. In light of the comparativist interpretation this means that there are gaps in an agent's comparative belief orderings. Given these two parts of the account, there are two components that need to be defended in order to defend a comparativist intersectionist interpretation. I will defend each of these components in turn. In this chapter I will defend the comparativist aspect of the account, and using this, in Chapter 3 I will defend the comparativist/intersectionist interpretation of credal sets.

Section 2.3 will look at the foundational question of what degrees of belief are, and using this, argue that to be an adequate account of degrees of belief, certain desiderata need to be met. In section 2.4 I outline the comparativist position and show how it is possible to represent an agent's comparative beliefs using a probability function. In section 2.5 I consider problems that have been presented for comparativism that suggest comparativism cannot meet the desiderata I have outlined. Going through these objections in turn I argue that comparativism is a viable account.

2.2 Intersectionism

Intersectionism is a particular interpretation of an imprecise credal set. In traditional Bayesian accounts we represent an agent's beliefs by a single probability function *Cr*.[1] Let Ω be a set of states of the world and $\mathcal{F}$ an algebra of sets on Ω and $Cr : \mathcal{F} \rightarrow \mathbb{R}$. The basic idea of imprecise credences is that rather than an agent's doxastic state being represented by a single probability function it can be represented by a set of probability functions. $C = \{Cr_1, Cr_2, Cr_3, \ldots\}$ where each of the Cr_i are probability functions and we take the set C to represent the agent's belief state. Broadly speaking an intersectionist interpretation takes the credal set as a whole to represent an agent's belief. That is, the credal set represents the agent's determinate attitude.

We can give an explanation for why we need an imprecise credal set to represent an agent's beliefs even when there is no indeterminacy in their belief by looking to comparativism. Comparativism holds that certain types of doxastic attitudes - comparative beliefs - play an important role in explaining numerical representations of an agent's at-

[1] Also referred to as a credence function.

titudes. Quantitative facts about belief can be reduced to an agent's comparative belief orderings. That is, claims of the form A is at least as likely as B. We can write this as $A \gtrsim B$. A probability function Cr represents an agent's comparative beliefs just in case

$$A \gtrsim B \text{ iff } Cr(A) \geq Cr(B)$$

If an agent's comparative beliefs are complete (and meet certain other conditions) then we can represent her beliefs with a single probability function. Completeness is the claim that for every all propositions A, B an agent has beliefs about it is either the case that A is at least as likely as B or B is at least as likely as A, or both.

However, this is seen as an unrealistically strong constraint on an agent's comparative beliefs by many comparativists. It seems likely that agents have incomplete comparative orderings. One might simply lack a comparative belief between two propositions. This can be represented if we take an imprecise credal set C to represent the agent's comparative beliefs rather than a single probability function.

With an imprecise credal set we can still represent claims about an agent's comparative beliefs:

$$A \gtrsim B \text{ iff } Cr_i(A) \geq Cr_i(B)\ \forall Cr_i \in C,$$

With an imprecise credal set we can also represent than an agent does not have certain comparative beliefs. We can represent that an agent fails to have any comparative beliefs between A and B. That is, we can represent that neither $A \gtrsim B$ nor $B \gtrsim A$ by an imprecise credal set C:

$$C = \{Cr_1, Cr_2, \ldots, Cr_n\} \text{ such that } Cr_1(A) > Cr_1(B) \text{ and } Cr_2(B) > Cr_2(A).$$

Comparativist intersectionism is therefore an interpretation of an imprecise credal set that allows us to make sense of an agent having an incomplete comparative ordering and thus an incomplete probability ranking.

2.3 Degrees of Belief

As noted above, we can give an interpretation of an imprecise credal set by appealing to a comparativist picture of degrees of belief. In particular, by taking agents to have incomplete comparative belief orderings. Comparativism is a view one might take on the nature of degrees of belief - it takes the position that comparative beliefs are more fundamental than partial beliefs and that from comparative beliefs we can explain how numerical degrees of beliefs represent an agent's beliefs. Given this, an argument for comparativist intersectionism depends on comparativism being a viable account of degrees of belief. In this chapter, I will give arguments for comparativism.

In order to judge whether the comparativist account is viable we need desiderata that an account of degrees of belief ought to meet. By looking at a historic account of degrees of belief in section 2.3.1 - the original betting interpretation, we can see areas where this account is widely deemed inadequate. Using this I will give desiderata for what an account of degrees of belief ought to look like in section 2.3.2. Using these desiderata I will argue that comparativism can meet them and thus provide us with an adequate account.

2.3.1 What are degrees of belief

When considering an agent's doxastic state there is the traditional framework of belief where an agent either believes or disbelieves or suspends judgment in a proposition. In contrast, we can also think of an agent's beliefs in a more fine-grained way with degrees of belief (which are also referred to as degrees of confidence or partial beliefs). An orthodox Bayesian view of degrees of belief is to think of them as having a numerical representation where that representation is a function that satisfies the probability calculus (for a rational agent at least). When an agent is fully confident in a proposition (i.e. in a tautology) their degree of belief is 1 and when an agent is fully unconfident in a proposition (i.e. in a contradiction) their degree of belief is 0.

There is a question about the relation between the doxastic state of having degrees of belief and the numerical representation of degrees of belief. It is unrealistic to think our beliefs are actually numerical. We have beliefs of the form 'I'm more confident it will rain in Leeds tomorrow than I am it will rain in Tehran' (comparative beliefs) and beliefs of the form 'I'm very confident that it will rain in Leeds tomorrow' (partial beliefs). We can also give different degrees of strength to our partial beliefs - 'I'm very confident it will rain in Leeds tomorrow' and 'I'm slightly confident we won't have another lockdown'. It seems unrealistic that with our different degrees of confidence in propositions we literally attach numbers to them in our head.

> the problem some nonprobabilists have with the notion of subjective probabilities . . . [is]that they find it difficult to conceive seriously and realistically of people as having "num-

> bers in the head" … (Zynda, 2000, p.50)

This idea is found elsewhere in the literature on comparativism with Stefánsson (2017) saying that "Numerical degrees of belief just ain't in the head" (p.8) and further noting that it is not just the case that we do not have numbers in our head but that we do not, or:

> … *should* not believe … that rational degrees of belief *are* probabilities … a probability function that represents a person's beliefs is not psychologically real. (Stefánsson, 2017, p.1)

If numerical degrees of belief aren't literally in our heads then we have to explain where they come from and then how they come to represent our beliefs.

Given these considerations, we want to give an account of what degrees of belief are that rests on something more fundamental and provides an explanation of how we can represent our degrees of belief with a probability function. By looking at the betting interpretation—an existing account of degrees of belief that is widely viewed as inadequate—we can identify certain features of what an adequate account ought to look like by looking at the shortcomings.

Betting Interpretation

The betting interpretation of degrees of belief is a special case of a more general class of theories about degrees of belief which tie them conceptually to an agent's (actual/dispositional/hypothetical) choices. The betting interpretation directly treats degrees of belief as an agent's betting dispositions/behaviour. de Finetti (2017) gives an account of degrees of belief (or probability) in terms of what an agent judges they would buy or sell a bet for. This gives us a means of being able to measure subjective probability.

> The probability $P(E)$ that You attribute to an event E is, therefore, the certain gain p that You judge equivalent to a unit gain conditional on the occurrence of E: in order to express it in a dimensionally correct way, it is preferable to take pS equivalent to S conditional on E, where S is any amount whatsoever, one Lira or one million, \$20 or £75. Since the possible values for a possible event E satisfy inf E = 0 and sup E = 1, for such an event we have $0 \leq P(E) \leq 1$, while necessarily $P(E) = 0$ for the impossible event, and $P(E) = 1$ for the certain event (de Finetti, 2017, p.64)

An attraction to identifying degrees of belief as not just behaviour but specifically betting behaviour and monetary amounts is it has the advantage of making degrees of belief sometimes observable and explaining their numerical nature in terms of this. The amount of money an agent would bet also makes it easy to explain where the numbers come from. With the betting interpretation, we also get the Dutch Book argument which justifies why an agent's degrees of belief ought to be probabilistic; if an agent's beliefs are not probabilistic then the agent will make a loss in a series of bets.[2] We assume that a rational agent would not want to make a loss and thus non-probabilistic degrees of belief are irrational. As such, some positive features of the bet-

[2] See for example de Finetti (1992) and Hájek (2009). Vineberg (2016) gives an overview.

ting interpretation are that it gives an explanation of why an agent's degrees of belief ought to obey the probability axioms and how beliefs can be used to inform (some) behaviour.

While it might be plausible that an agent's betting behaviour can sometimes convey their degree of confidence in an event, it is entirely unrealistic to think that an agent's betting behaviour is really all there is to degrees of belief. There are many well-known objections to the betting interpretation in the form presented above. If degrees of belief were just betting dispositions then it should be possible to increase the value of the bet so long as the ratio stays the same. However, it seems unlikely that this would be the case based on the wealth of the agent making the bet. An agent might be willing to bet $50p$ to $£1$ but not willing to bet $£50,000$ to $£100,000$. The betting interpretation also doesn't account for risk-adverse behaviour or the possibility that an agent could have a strong aversion to betting (and hence possibly refuse all bets even though they still have beliefs).[3]

It seems the above problems stem from saying degrees of belief are literally an agent's betting behaviour. This suggests that to avoid these problems we might need to look beyond the actual bets an agent makes. We could instead consider an agent's betting dispositions and identify an agent's degrees of belief with the price an agent would be disposed to bet. This raises the question of what it means to be disposed to bet in a certain way. We could, for example, think about an agent's dispositions to bet in close possible worlds. However, if

[3] In general the objections to behaviourism apply to this account.

we are only considering close possible worlds the same objections as in the actual case will apply. In close possible worlds, we should expect the agent to not be willing to make scaled bets for vast sums of money and it is just as possible that in a close possible world an agent has a credence in a proposition but no disposition to bet in a certain way. We could instead look to idealised possible worlds where the agent always has the corresponding dispositions to bet. However, these worlds are so far removed from the actual world that it is no longer clear how they can factor into an explanation of what an agent's degrees of beliefs in this world are. It seems then that turning to betting dispositions is still inadequate as an explanation of what degrees of belief are.[4]

In the betting interpretation, where degrees of belief are tied to an agent's betting dispositions, we can see how this only seems to work under specific assumptions. The agent is not unwilling to gamble on principle, the agent is not risk-averse, the value of the bet is not so high that they reject the bet on those grounds etc. The betting interpretation is one of a number of accounts that links an agent's preferences, behaviour and degrees of belief.[5] There are also accounts that show how degrees of belief are related to preferences via a decision

[4]See Eriksson and Hájek (2007) for a more in-depth version of this argument.

[5]I take preferences and choice behaviour to be closely connected. We can often infer choice behaviour from an agent's preferences and vice versa, "Preference is linked to hypothetical choice, and choice to revealed preference."(Hansson and Grüne-Yanoff, 2021, §5) See also (Savage, 1954, p.17) who expresses a view about how preferences are related to decisions to act (preferences should (at least in principle) be determined by decisions to act.)

theory based representation theorem.[6] The rough idea behind representation theorems is that we can use a probability function (credences) and utility function to rank options, if an agent's preference ranking conforms to certain axioms we can represent their preferences using expected utility maximisation in relation to these credences and utilities.[7]

Eriksson and Hájek (2007) note that there is a general worry we should have with any account that posits such a direct link between an agent's preferences and their degrees of beliefs. Namely, we cannot identify credences with preferences since it is possible for them to come apart. Accounts that attempt to derive degrees of belief from preferences rely on an agent's preferences satisfying some very strong conditions. Just as it seems unrealistic and unintuitive that an agent satisfies all the conditions needed for the betting interpretation to work, we might also think that a number of the conditions needed for preference-based accounts are unintuitive. This can be illustrated with the example of the Zen Monk who has credences but no preferences.

> Gazing peacefully at the scene before him, he believes that Mt. Everest stands at the other side of the valley, that K2 does not, and so on. But don't ask him to bet on these propositions, for he is indifferent among all things. If the

[6]See Ramsey (1926), Savage (1954), Anscombe, Aumann, et al. (1963), Cozic and Hill (2015).

[7]There are also functionalist and interpretivist accounts that also have preferences play a central role such as Maher (1993) and Lewis (1974).

> monk is conceptually possible, then any account that conceptually ties credences to preferences is refuted. (Eriksson and Hájek, 2007, p.194)

The general worry is that all these preference-based or choice-based accounts can only reliably tie credences to preferences/choices given very special conditions which aren't intuitively necessary for having degrees of belief, and moreover (in some cases) are not very realistic. The Zen Monk does seem conceptually possible and so it seems we cannot use the betting interpretation to provide an understanding of how numerical degrees of belief come to have representational content. This suggests then that an account of what degrees of belief are should not identify degrees of belief with preferences alone.[8]

2.3.2 Desiderata for an account

In order to assess whether an account of what degrees of belief are is adequate we need to give certain criteria that an adequate should fulfil. The problems with the betting interpretation demonstrate that an account of what degrees of belief are should be more psychologically realistic and not only link an agent's beliefs to their actual behaviour, behavioural dispositions or preferences. It should be possible to use the account to explain why we think the beliefs of non-ideal agents can be represented by numerical degrees of belief.

In this section I will consider the key roles that degrees of belief play

---

[8] This is not to say that all preference-based accounts face this problem. As Elliott (2019) notes there are a range of preference-based accounts and more complex accounts of how preferences and numerical degrees of belief relate may not be concerned with the Zen Monk argument.

and from these give desiderata for an account. When looking at an account of belief we want it to be able to explain the functional role of degrees of belief as it appears in economics, folk psychology and our best normative theories. That is to say, we want a theory that fits with and explains the role degrees of belief play in explaining an agent's beliefs at a certain time, how an agent updates her beliefs (or how she ought to update her beliefs via conditionalization) and the role of degrees of belief in an agent's decision making, i.e the role of an agent's beliefs in making decisions in both decision theory and in game theory. Given this we see that for an account of degrees of belief to fit these roles we want the following desiderata to be met:

(i) **Possible for an ordinary human agent to have degrees of beliefs.**

It should be possible to use the account to explain why we think the beliefs of non-ideal agents can be represented by numerical degrees of belief. Relatedly, the account should be able to explain our folk psychological notions of belief or at least " allow for folk psychological explanations of people's choices" (Stefánsson, 2017, p.383). An account that can make sense of our everyday notions of belief, rationalising and choices is desirable for a number of reasons. For one, we can use talk about partial beliefs in everyday language without any associated technicalities such as a theory of probability or decision theory. I might say I'm very confident in *A* but only slightly confident in *B*.[9]

[9] Eriksson and Hájek (2007) argue that we already have mastery of the concept of degree of belief and that formal theories of probability as applied to this concept

(ii) **An agent's degrees of belief can be represented by probabilities (in some sense).**

There are a number of arguments that have been given that show why rational agents beliefs ought to conform to the probability calculus and in particular to be probabilistically coherent. As I mentioned in Chapter 1 there are arguments for probabilistic coherence from accuracy considerations and dutch book arguments. Representation by a precise probability function may be too strong a requirement, however. We may wish to represent that an agent has incomplete beliefs. In this case rather than a representation by a single probability function we give an imprecise probability representation.

(iii) **We can explain cardinal facts about belief.**

From an agent's comparative beliefs we can easily read off information of the form 'an agent thinks *A* is more likely than *B*'. However, it seems like we also have beliefs that require cardinal information such as an agent believes *A* is much more likely than *B* or even an agent thinks *A* is twice as likely as *B*. Since it seems that we do have beliefs of this form an account of degrees of belief should be able to accommodate and explain cardinal facts about belief. Moreover, it seems that we can observe that people behave as if they have cardinal facts about belief (i.e I might always make bets as if I think *A* is twice as likely as *B* - cardinal

only add structure to a concept that we already have in our folk psychology.

facts about belief can explain why I behave in this way).[10]

(iv) **Be able to account for interpersonal comparisons of beliefs (i.e being able to compare Ann and Bob's beliefs).**

We want to be able to explain why we think Ann is more confident it will rain than Bob is (Stefánsson, 2018, p.383). Interpersonal comparisons of belief allow us to make sense of differences in people's behaviour. For example, Ann may pick up an umbrella before she leaves the house, but Bob does not. We can explain this difference in behaviour (in part) by comparing Ann and Bob's confidence in the likelihood of it raining.

(v) **Be able to account for intrapersonal comparisons between your own beliefs at a time t_1 and a time t_2.**

It seems like an agent should be able to express that they are more confident in A at time t_2 than she was at time t_1. This also seems necessary if we want to be able to explain differences in an agent's behaviour at different times. For example, if Andi is twice as confident in A than B at time t_1 but only slightly more confident in A than B at time t_2 and this difference is reflected in her behaviour at these different times.

In addition, a desirable feature of an account is that it does not require an agent have precise credences. As discussed above, precise credences seem like an unrealistically strong constraint for real agents.

[10] A criticism of comparativism has been given by Meacham and Weisberg (2011) who claim that if we accept comparativism we have to sacrifice all cardinal information about belief.

If we accept precise credence functions then we also accept that for any two propositions *A* and *B* an agent has beliefs about, one of the following relations must hold: the agent thinks *A* is more likely than *B*, the agent thinks *B* is more likely than *A*, or the agent thinks *A* and *B* are equally likely; i.e. precise credence functions require completeness of opinion. This seems like an unrealistically strong constraint. Often an agent simply doesn't have an opinion. Any account that allows for incomplete comparative orderings is more psychologically realistic than one that does not.

As well as being more psychologically plausible for real agents a lack an opinion it also seems like the desirable view to adopt for ideally rational agents in some cases, i.e. in cases where there is a lack of sufficient evidence for an agent to justifiably form a precise credence. Despite lack of evidence the precise credal accounts require agents to have rich credal sets that are complete. Another reason to favour accounts that accommodate imprecise credences is for information theoretic reasons - those who favour maximum entropy methods will argue that we want to have the least informative doxastic state consistent with the evidence.[11] An incomplete comparative belief relation will be less informative than complete belief relations and so, if the aim is to minimise informativeness we will often be required to opt for the incomplete comparative belief relation (see (Konek, 2019, p.274)).

[11] See Abellan and Moral (2003) where they talk about measures for maximum entropy for credal sets. Konek (2019) suggests that the entropy measures they discuss might also serve as measures for comparative beliefs.

2.4 Comparativism

Comparativism is the view that comparative belief relations are more fundamental than partial beliefs. Comparative belief is taken to be a fundamental belief relation and from this, we can give an explanation of how we can arrive at numerical degrees of belief (and therefore how numerical degrees of belief can represent an agent's beliefs).

> we do not believe ... that rational degrees of belief *are* probabilities. We do ... believe, however, that rational beliefs have a structure that ensures that they can be *represented by* a (perhaps non-unique) probability function. But a probability function that represents a person's beliefs is not psychologically real. Nor are the axioms of probability theory normatively fundamental. What is psychologically real are comparative belief relations, and what is normatively fundamental are requirements on the structure of these relations (Stefánsson, 2017, p.1).

There are a number of reasons why we should think comparative belief relations are important in explaining and modelling an agent's beliefs. We clearly do have comparative beliefs of the form 'I think A is more likely than B'. It also seems plausible that we can have these comparative beliefs without also having associated numerical degrees of belief about A and B. On the other hand, it also looks like if an agent has partial beliefs (degrees of belief) that they must also have comparative beliefs, if I have high confidence in A and low confidence in B then it follows that I think A is more likely than B. This implies that comparative belief relations are more fundamental than degrees of belief.[12]

As noted, regarding desiderata for an account, it is desirable that an account of what degrees of belief are does not constrain us to precise credences and can instead account for incomplete comparative orderings.[13] When we consider an agent's comparative beliefs we can see that often real agents don't have complete opinions. Konek (2019) gives the following example:

> Let C be the propositions: Copper will be greater than £2/lb in 2025.
>
> Let N be the proposition: Nickel will be greater than £3/lb in 2025.

Konek says "I am not more confident of ...C than ...N. Neither am I *less* confident in C than N, nor *equally* confident. I simply lack an opinion on the matter" (p. 273). We can model this lack of opinion with comparative beliefs by simply choosing a comparative belief relation $\gtrsim$ such that:

$$C \not\gtrsim N \text{ and } N \not\gtrsim C$$

In the next section I will show that comparative beliefs are more fundamental than credences and that we can use comparative beliefs as

[12] Thanks to Ed Elliott here for discussion.

[13] Other accounts of degrees of belief also give us imprecise credence models (such as preference-based accounts where agents have incomplete preferences orderings). So this feature of comparativism is not unique, but it is good that as an account it makes sense to say that we do not have complete comparative belief orderings.

part of an explanation of what credences are, but we cannot use credences as part of an explanation of what comparative belief relations are. This provides the strongest reason for taking comparative beliefs as more fundamental. It also shows that comparativism does more than measure numerical degrees of belief. Rather, it explains how we can give a numerical representation of belief given an agent's comparative beliefs.

2.4.1 Comparative Probability

The basic structure of an explanation for how our beliefs can be represented by probabilities is via results that show that if comparative beliefs have certain formal structures they can be represented by a probability function. We then claim that the comparative beliefs do in fact have this formal structure, or at least approximate it, and thus they are in fact representable by probability functions.

There are a number of accounts of comparative probability and how we can get a probabilistic structure. Let us consider the comparative belief relation $\succsim$. Where the interpretation of $A \succsim B$ is that A is at least as probable as B. $A \succ B$ is interpreted as A is strictly more probable than B (and defined as $A \succ B$ if and only if $(A \succsim B)$ & $\neg(B \succsim A)$), $A \sim B$ is interpreted as equally probable and defined as $A \sim B$ if and only if $((A \succsim B)$ & $(B \succsim A))$. Below I will outline a weak requirement on the comparative belief relation that is a necessary condition for the relation $\succsim$ to be a probability. Recall the definition of a probability function.

Definition 2.1. Let Ω be a set of states of the world and $\mathcal{F}$ a Boolean algebra of sets on Ω (an algebra of sets must be closed under complemetation and unions and intersections). Let Cr be a function from $\mathcal{F}$ to the real numbers. Cr is a probability function if and only if for every $A, B \in \mathcal{F}$:

1. $Cr(A) \geq 0$
2. $Cr(\Omega) = 1$
3. If $A \cap B = \emptyset$ then $Cr(A \cup B) = Cr(A) + Cr(B)$

Let $\succsim$ denote a comparative belief relation defined on $\mathcal{F}$. A function f represents $\succsim$ just in case for any $A, B \in \mathcal{F}$, $f(A) \geq f(B)$ if and only if $A \succsim B$. Note this function need not be a probability function. A necessary condition for a relation to be representable by a probability function is that it be a qualitative probability (Krantz et al., 1971, §5.2.1).

Definition 2.2. [Qualitative Probability] The relation $\succsim$ on $\mathcal{F}$ is a qualitative probability if and only if for any $A, B, C, D \in \mathcal{F}$ it satisfies the following:

A1. Weak Ordering:

  (a) Transitivity: if $A \succsim B$ and $B \succsim C$ then $A \succsim C$.

  (b) Complete: either $A \succsim B$ or $B \succsim A$

A2. Normality:

  (a) $\Omega \succ \emptyset$

  (b) $A \succsim \emptyset$

A3. Qualitative Additivity:
If $A \cap B = A \cap C = \emptyset$, then $B \succsim C$ iff $(A \cup B) \succsim (A \cup C)$.

Is it plausible that an agent has comparative belief relations that meet these necessary conditions to be represented by a probability function? Looking at the necessary conditions most of these conditions look like plausible requirements on comparative belief. A2. says that Ω is strictly more probable than the empty set (i.e. a contradiction) and that every proposition is at least as probable as a contradiction which is exactly what we want a comparative belief to say. The requirement of transitivity in A1. looks like a natural requirement on comparative belief. A3. says that if there is some *A* that is incompatible with both *B* and *C* then disjoining *A* to *B* and *A* to *C* does not affect the comparative probability of *B* and *C*. While these conditions just mentioned look like plausible requirements there is some contention about the completeness condition in A1.[14] Hawthorne (2016) and Alon and Lehrer (2014) solve this by using credal sets, it is possible to represent an incomplete comparative ordering if we take a set of probability functions.

These give necessary but not sufficient conditions. Both Kraft et al. (1959) and Scott (1964) show that we can give stronger axioms that are necessary and sufficient conditions for a relation to be probabilistically representable. I will focus on Scott's axioms.[15]

Theorem 2.1. Let $\mathcal{F}$ be a finite Boolean algebra and let $\succsim$ be a binary relation on $\mathcal{F}$. For $\succsim$ to be realisable by a probability measure on $\mathcal{F}$ it is necessary and sufficient that the conditions:

1. Non-Triviality: $\Omega \succ \emptyset$

[14] See (Fishburn, 1986, p.339).

[15] (Scott, 1964, p.246). Presentation of Scott's axioms has been take from (Konek, 2019, p.277).

2. Non-Negativity: $A \gtrsim \emptyset$

3. Complete: either $A \gtrsim B$ or $B \gtrsim A$

4. Isovalence: If $A_1 + A_2 + \cdots + A_n = B_1 + B_2 + \cdots + B_n$ and $A_i \gtrsim B_i \; \forall i \leq n$ then $A_i \lesssim B_i \; \forall i \leq n$ as well.

The axiom of note that differs from those given for a qualitative probability is axiom 4. As Konek puts it, if $\{A_1, A_2 \ldots A_n\}$ and $\{B_1, B_2 \ldots B_n\}$ are two sets of propositions that are isovalent, that is to say, they contain the same number of truths in every possible world. Then axiom 4 says "you cannot think that the A_is are uniformly more plausible than the B_is" (p. 277). Where, by uniformly more plausible he means that for each i you think A_i is at least as plausible as B_i and for some j there is an A_j strictly more plausible than B_i.[16]

2.4.2 Ordinal vs Cardinal

An obvious worry with the comparative belief framework is that it provides us with ordinal information about an agent's belief rather than cardinal. Ordinal information is of the form Andi believes A more than she believes B. Whereas cardinal information is of the form Andi believes A twice as much as she believes B. Elliott gives the following example of an agent Sally about to role an ordinary 6 sided die:

> ORDINAL. Sally is *more* confident of rolling ≥ 2 than of rolling a 1.
>
> INTERVAL. Sally is *much more* confident of rolling ≥ 2 than

[16] There are a number of ways to ensure a unique probabilistic representation. See Fishburn (1986) and (Kraft et al., 1959, §5.2).

of rolling a 1

RATIO. Sally is *five times* as confident of rolling ≥ 2 than of rolling a 1. (Elliott, 2020b, p.5)

As noted in my desiderata for a concept of belief we want to be able to talk about the cardinal information of a belief. However, it looks like it may be problematic for a comparativist to give either interval or ratio information about belief. Meacham and Weisberg (2011) have the pessimistic view that if we accept comparativism we also accept that we have no cardinal information saying "one must be prepared to sacrifice all cardinal facts about degrees of belief: absolute values, differences, and even ratios of differences, are all unreal on such a view" (p.658). As I will discuss in more detail in section 2.5 one worry is that if we grant that there are numerical representations of comparative beliefs we can give non-probabilistic representations. What is 'real' about these different representations is what they share (in virtue of the underlying comparative belief ordering) then we might worry that there is no common interval or ratio information across these numerical representations.

There is however a standard explanation by comparativists of how we can get cardinal facts about belief from the ordinal information about belief given by comparative belief relations. The explanation is via an analogue to measurement or length or weight. Fine gives an example by considering objects of different weights (see (T. Fine, 1973, p.68)). I will use a related example by considering two concrete objects α and β which we have ordinal information about (namely that α is longer than β) and show how we can give the conditions under which we can say that α is twice as long as β (without presupposing cardinal information).[17]

The trick is to suppose that there is another, distinct, concrete object which we will call β^* which is the same length as β. The length comparison of β and β^* can therefore be done just with ordinal information. If we also have the resources to join end-to-end the two objects β and β^* then we have all we need to be able to say that α is twice as long as β:

> α is twice as long as β if and only if α is the same length as the composite object of β and β^* joined together end-to-end.

Similarly, we can extend the same method to three objects to be able to say what it means for λ to be three times as long as β. If there are two concrete objects β^* and β' which are both distinct from β and distinct from each other and are all the same length then, by applying our method of concatenation again we can say that λ is three times as long as β if and only if λ is the same length as the composite object of β and β^* and β' joined end-to-end-to-end.

We can thus see how it is possible, in the case of length, to get ratio information. That is, we can give a story for how to give numerical lengths to objects that encode ratio information. As Elliott (2020b) notes there are some assumptions that need to be made in order for this example to work and give the right truth conditions for α being twice as long as β. For example, we need there to be enough objects in the world that we can generate all the ratio information - given the

[17] See Elliott (2020a,b).

number of objects in the world, it is unlikely that this will be a problem for the case of measuring concrete objects but the analogous requirement on beliefs (that there be enough beliefs so we can generate all the ratio information may be more problematic). We also need the *at least as long as* relation that exists between objects that satisfies the requisite axioms.

Mapping this over to comparative beliefs we can see that there are certain assumptions that we will have to make if this sort of approach is to be plausibly used to give us cardinal information about belief. We need to give some operation that is analogous to the concatenation of two concrete objects then we can concatenate two beliefs and use this to explain that an agent is twice as confident in A as B. That is, we need some addition operator. An addition operator seems unproblematic for the comparativist however, since, as we have seen in the previous section, we can represent the agent's beliefs by a probability function. We can use the union of disjoint propositions as our addition operation. As we can see from the axioms, if we have two disjoint propositions A and B then the probability of the union of A and B is the probability of A added to the probability of B. The union of disjoint propositions, therefore, looks like an ideal concatenation operation. Indeed, it forms the basis of Stefánsson's Ratio Principle which he uses to explain how we can get ratio information from comparative probabilities:

> **Ratio Principle** (RP). We say that an agent, with a continuous qualitative probability relation $\succsim$ on Ω, is twice as confident in A as in B, just in case there is a $B^* \in \Omega$ such that: (i) $B \sim B^*$, (ii) $B \cap B^* = \emptyset$, (iii) $B \cup B^* \sim A$.[18] (Stefánsson, 2018, p.385)

The union of disjoint propositions is also used by (T. Fine, 1973, p.68) (Krantz et al., 1971, p.200) and DiBella (2018) in their proposals of how we can concatenate beliefs. We can therefore see how it is possible to get ratio information from ordinal information.

This gives us a story of how we can give a probabilistic representation of beliefs. It does require that an agent's comparative beliefs have a certain structure in order to be representable by a probability function. There are also requirements on the number of beliefs an agent has (so we can give confidence duplicates). I have pointed towards some reasons why these seem like reasonable assumptions to have about an agent's comparative beliefs, however, it does seem unrealistic that real agents will meet all these requirements. This opens up some potential objections to the account. In the next section I will outline some prominent problems that have been raised for comparativism and argue that these problems can be resolved.

[18] Notation altered for consistency.

2.5 Problems for Comparativism

There are a number of problems that have been presented for comparativism. These problems highlight difficulties the comparativist account may have in meeting the desiderata I outlined above including explaining the role of numerical beliefs in decision theory, explaining interpersonal beliefs and whether the conditions required for comparative beliefs to be probabilistically representable are too strong. In this section I will outline examples that show how these problems might arise and address how the comparativist can account for these cases.

In sections 2.5.1 and 2.5.2 I will consider problems that arise for comparativism in light of the conditions required for a comparative belief ordering to be representable by a probability function. In section 2.5.1 I look at the case where an agent's comparative beliefs fail to meet the conditions required for a unique probabilistic representation. Multiple probabilistic representations of the same comparative beliefs pose a problem for this accounts ability to explain how degrees of belief integrate with decision theory. I argue that this can be accommodated by the comparativist theory by taking the multitude of probabilistic representations to collectively be a representation of an agent's comparative beliefs.

In section 2.5.2 I consider the possibility of an agent's comparative beliefs failing to meet the conditions required for *any* probabilistic representation. The requirements for probabilistic representation seem too strong for real agents to meet. I will argue that so long as an agent's comparative beliefs *approximately* meet the conditions for probabilistic representation we can give a probabilistic representation of an agent's comparative beliefs.

In section 2.5.3 I will consider whether comparativism can accommodate interpersonal facts about belief.

2.5.1 Multiple Probabilistic Representations

One issue for comparativism is that comparative beliefs seem too coarse-grained to give a unique representation. There are two ways this can come about: there may be multiple probability measures that can represent the same comparative beliefs and there will always be other numerical representations as well as probabilities that can represent the comparative beliefs. A general worry for non-uniqueness is whether the existence of other representations raises the question of whether "*being representable as having* certain degrees of belief, as described by a probability function p, is sufficient for *really having* those degrees of belief." (Zynda, 2000, p. 49). If there are multiple probability functions compatible with an agent's comparative beliefs there is no reason to think one representation is privileged over any other, there is no 'real' representation. Taking each of these ways of non-uniqueness in turn I will illustrate in more detail how the non-uniqueness may come about and how we can respond to them.

If there are multiple probability representations of the same comparative ordering this looks problematic for the role of degrees of belief in decision theory. We can give a simple example of an agent's beliefs at two different times where her comparative belief ordering stays the same but intuitively it looks like we can ascribe different partial beliefs to her.

Consider an agent Andi and her beliefs at times t_1 and t_2. She has comparative beliefs with respect to three propositions A_1,A_2 and A_3 and her beliefs are probabilistically representable. Her comparative beliefs at both t_1 and t_2 are the same and the following holds of her beliefs: $(A_1 \cup A_2) \sim A_3 > A_2 > A_1$. We can see that a number of different probability functions can represent this one comparative belief ordering. Indeed, it is numerically possible and also conceivable that she has different partial beliefs towards A_1,A_2 and A_3. We can imagine that at t_1 she is more than twice as confident in A_2 as A_1, and that she changes her belief at t_2 so her comparative belief ranking stays the same but she is now only a little more confident in A_2 than she is in A_1.[19]

The above situation seems conceivable. Indeed, given the coarse-grained nature of the comparative beliefs Andi has, it is easy to see how multiple probability functions can be consistent with the comparative belief rankings. However, it seems like there should be some important and meaningful distinction to be made here. The probabilistic comparativist cannot accept this. Since Andi has the same comparative beliefs, and since all there is to belief is comparative beliefs, then she has the same beliefs at the different times. Even though it looks like we can conceive of a difference between her beliefs at t_1 and t_2 this difference is not real. While this looks problematic one

[19] Example can be found in (Elliott, 2020b, p.13).

might simply bite the bullet here. The situation may seem conceivable but just because it is conceivable does not mean it will actually occur.

Having said that, Elliott shows that we can move beyond it merely being a conceivable difference. We can consider decision situations where the difference matters such as the following given in (Elliott, 2020b, p.14):

	A_1	A_2	A_3
Option α	$-2x$	x	x
Option β	0	0	x

Where, if Andi had the beliefs described above at t_1 and t_2 she will choose option α at time t_1 and option β at time t_2 even though her comparative belief ordering stays the same. It seems then we can give decision situations where Andi has different preferences (and different betting behaviour) but the comparativist cannot say this is down to any difference in belief because there is none.

This example serves to illustrate the general point that a probability function gives us exactly one belief ranking but this does not hold in the other direction. It is possible that belief rankings can be represented by multiple probability functions where the difference in those probability functions can conceivably lead the agent to make a different choice in decision scenarios such as in betting scenarios. If an agent consistently makes choices that seem to indicate that they have more fine-grained degrees of belief than the comparative ordering shows then this would show that comparativism cannot explain an agent's degrees of belief adequately. There are conditions under which there will be a unique probability function for a comparative belief ranking, but in many cases, we will be able to give examples like the one above where the comparativist ought to give an explanation. As noted in the desiderata, an adequate account of degrees of belief ought to be able to account for the functional role degrees of belief play in decision theory. So, we need an explanation about the apparent difference in degrees of belief that arises from decision situations Andi might be put in.

Looking at the example above it is quite contrived in the sense that Andi has very few comparative beliefs. So we can give examples where it looks like she can have extreme difference between her confidence in A_1 and A_2 and it not change her comparative belief ordering. In this example, it is unclear to me why we should have any intuitions about there being different cardinal information about Andi's beliefs. It is unclear why we should ascribe more structure to her beliefs than the comparative belief ordering. It seems that the intuitions about it being possible for her to have these differences in her beliefs come from similar situations where she has a richer comparative belief ordering and thus we can make sense of her beliefs changing in terms of her comparative belief ordering changing. If we take richer comparative belief orderings then we will not get such extreme versions of this example where there can be such large differences in the bets she will accept at different times and yet she has the same comparative ordering. There is a disconnect between the intuitions that Andi can have these differences in her credences that cannot be explained by the comparative ordering because Andi does not have enough confidence-duplicates to determine ratios for all pairs of propositions. Given enough confidence-duplicates, we can explain what it means for her to have a change in confidence between times t_1 and t_2 (and there will be a change in comparative belief ordering as well). Once we remove such extreme examples then it may be that our intuitions about what we can say in response changes.

Even so, one might respond that the challenge will remain. Often a comparative belief ordering is not fine-grained enough to give a unique probability ordering and we can come up with a decision situation that will appear to show an agent has important differences in her confidence levels at different times but with no change in her comparative belief ordering.

There are two cases to consider here: the ideal agent and the non-ideal agent. Considering the non-ideal agent, it looks like this is illustrating a real feature of agents beliefs, namely that we do have coarse-grained beliefs and that betting behaviour may in some sense be an arbitrary choice by agents. I take the coarse-grained nature of comparative beliefs to show that we typically have vague or imprecise attitudes (not numerically precise ones that are representable by a unique probability function). So, a solution to the problem of multiple representability by probability functions is to take all the possible probability representations as a whole to be representational and therefore treat the agent's credence as vague or imprecise. Each probability function can be understood as a credence function in the agent's credal set.

For an ideal agent, one might argue that we cannot appeal to it being more psychologically realistic for an agent to have coarse-grained comparative beliefs. This may be, but if in the ideal agent case we can appeal to unrealistic detail then it looks like there are again two cases to consider. One case is where an ideal agent has comparative beliefs that meet the requirements for them to be represented by a unique probability function. In this case, we don't have a problem, comparativism can explain where the degrees of belief come from and their role in decision theory. The second case to consider is one where an ideal agent doesn't have a unique probability ordering. In this case, we can once again argue that if the agent has coarse-grained beliefs their betting behaviour may be an arbitrary choice. It is unclear to me why we ought to have intuitions that an ideal agent could simultaneously have coarse-grained comparative beliefs but fine-grained preferences.

Fine (1973) also shows how we can give a decision theory in comparative probability framework (although he notes that there is much more work to be done in developing the account further). Since an agent's comparative beliefs may be coarse-grained and therefore multiple probability functions may represent their comparative belief ordering it suggests that we should view their decision making in comparativist decision theory. On Fine's approach we view chance and uncertainty comparatively rather than quantitatively.[20] This shows that we can recapture elements of decision theory in purely comparative terms as opposed to probabilistic terms.

[20] See also T. Fine (1971).

2.5.2 No Probabilistic Representation

Another substantial worry for the comparativist view is the possibility of agents who have comparative belief orderings that aren't representable by any probability function (or set of probability functions). There are several requirements in the above story for how we can give a probabilistic representation of an agent's comparative beliefs. We might worry that some or many of these requirements are unrealistic or too hard for an agent to meet. In which case, even though we have a story for how it might be possible to give a probabilistic representation from an agent's comparative beliefs this story might not be helpful for non-idealised agents that do not have sufficient beliefs for confidence duplicates to give cardinal facts about belief.

As we can see, there are necessary and sufficient conditions on the structure of an agent's comparative beliefs in order for them to be representable by a probability function (or a set of probability functions). Looking at these requirements we can see that some of them might be unrealistic for an irrational agent to meet, and it seems like real human agents will be irrational. Indeed, there are a number of examples that seem to show that real agents are irrational such as the Ellsberg paradox[21] and the Conjunction Fallacy.

The conjunction fallacy shows that the condition of monotonicity (which I have referred to as qualitative probability in Definition 2.2) seems to

[21] In the Ellsberg paradox the agent is asked their preferences between various bets. The set-up is such that some of these bets are over ambiguous outcomes and we see people tend to have ambiguity aversion - that is they tend to prefer to bet on known rather than ambiguous probabilities. See Ellsberg (1961).

fail for real agents. If this is the case, then one of the conditions for a comparative belief ordering to be probabilistically representable fails and so it seems like we cannot represent an ordinary agent's comparative beliefs with a probability function. An example of the conjunction fallacy goes as follows, we're told:

> Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. (Kahneman et al., 1982, p.92)

People are then asked what they think is more probable.

A. Linda is a bank teller.

B. Linda is active in the feminist movement.

$A \cap B$. Linda is a bank teller and active in the feminist movement.

A significant number of people when asked will answer that they think $A \cap B$ is more probable than A. The probability of the conjunction of two events is always less than or equal to the probability of one of its conjuncts. That is to say $P(A \wedge B) \leq P(A)$ and $P(A \wedge B) \leq P(B)$. However, as we can see this gives non-monotonic comparative beliefs and so there are cases where an agent's beliefs fail to satisfy the requirements for the measurement analogy. Do these people still have degrees of belief?

From this, it follows that we can't give cardinal information about an agent who has these non-monotonic beliefs. However, as Elliott (2020b) points out it seems like, even in this example where monotonicity fails that there could still be cardinal information. That is, there is an intuition that one agent might think $A \cap B$ is much more likely than A while another might think $A \cap B$ is only slightly more likely than A. However, given the measurement analogy and the requirement of monotonicity we cannot say this.

What do we say about people whose comparative orderings don't satisfy the axioms required for probabilistic representation? In section 2.4.1 I outlined the necessary and sufficient conditions for a comparative ordering to be probabilistically representable. As I argued, these conditions mostly seem reasonable conditions to assume of an ideal rational agent (with the exception of the completeness condition but by taking imprecise credences we can remove this requirement). So, any problem of lack of probabilistic representability seems to be focused on irrational agents. From the desiderata, we can see that we want a theory of degrees of belief that can account for real agents beliefs and it looks like real agents will often have irrational beliefs. Indeed, something like the conjunction fallacy example is a case where an agent is being irrational in her beliefs.

When we look at these examples there are explanations about why it looks like lots of people make the same irrational judgements. Some of these explanations aim to show that people are in fact making rational judgements.[22] Or, at least the agent is not necessarily making

[22] In the case of the conjunction fallacy for example it has been argued by Hartmann and Meijs (2012) that it might not be irrational to judge $A \cap B$ as more likely than A. They give an example that is different but similar to the Linda problem - the Walter problem. In the Walter problem we are presented with similar information to the Linda case but unlike in the Linda case where we are told the information about her is certain in the Walter problem we are told by a person and the extent to

an irrational judgment.[23] So we might think that experiments that purport to show that real agent's systematically have non-monotonic beliefs or irrational beliefs do not in fact do so. However, it seems that often we do see that real agents are not probabilistically coherent.[24] It does seem that real agents can be closer or further away from being ideally rational.

When considering irrational agents we can consider degrees of irrationality. That is, some agents might have comparative beliefs that are very near to meeting the necessary and sufficient conditions to be probabilistically representable (let us refer to comparative belief orders that meet these conditions as 'ideal comparative orderings'). Other comparative belief orderings might be very far away from an ideal comparative ordering. Using this we can say that we accept real agents fall short of meeting all the conditions for ideal comparative orderings but so long as an agent is only slightly irrational we can approximate their comparative beliefs with a probability function.

A similar idea of degrees of irrationality can be found in Staffel (2019b).

---

which we trust the persons information contributes to what sorts of judgements one might make about Walter. They show that if we take a conjunction of propositions conditional on the evidence (that is the reports of the person telling us about Walter) then it can be the case that it is more probable than the conjuncts condition on the evidence.

[23] Wolford et al. (1990) suggest that it is not really a fallacy as people are making judgements of the form $P(X \mid A\&B) > P(X \mid A)$ which is not necessarily irrational rather than $P(A\&B \mid X) > P(A \mid X)$. They argue that in the Linda example it is not clear which rule to apply.

[24] Staffel (2019b) considers this one of the assumptions of Bayesianism that it seems that non-ideal agents can be more or less irrational (where an agent is irrational if they fail to meet normative requirements such as the norm of Probabilism). See also (Titelbaum, 2019, §4) who gives an overview of arguments against credal constraints (since certain credal constraints might seem unrealistic for real agents).

She considers agents that have probabilistically representable beliefs (credences) that do not meet the rational ideal of probabilistic coherence. It seems that agents can be closer or further away from the ideal of coherence. Consider two agents David and Arthur who both have credences in the same set of propositions. Ideal agents ought to have credence 1 in tautologies. There is one tautology T that neither of them has credence 1 in, they are therefore both irrational. However, Arthur has credence 0.99 in T whereas David has credence 0.5 in T. Intuitively it looks like even though they are both irrational, Arthur's credences are closer to an ideally coherent credence.[25] Staffel shows that we can formalise the idea of closeness or being a better approximation to probabilistic coherence by giving a distance measure. Then we can measure how far away from the closest coherent representation an agent's beliefs are.[26]

Taking this idea, we can give a distance measure for the difference between two comparative belief orderings. Bogart (1973) for example, gives a distance measure between transitive preference relations (the formal structure of which is the same as a transitive comparative belief relation). We can apply such a distance measure to comparative belief orderings and measure how close a comparative belief ordering is to its closest ideal comparative ordering. Using this we can say that if an agent David has a non-ideal comparative belief ordering that is within a certain distance γ from an ideal comparative ordering

[25]Indeed the idea that real agents are only approximate Bayesian agents has been argued for in a number of places (see Griffiths et al. (2012), Chater et al. (2011)).

[26]See (Staffel, 2019b, Ch.3).

which can be represented by a probability function Cr_d then, we can say David has degrees of belief that are approximately represented by Cr_d. This can explain how, in general, an agent can have degrees of belief on the comparativist picture without meeting all the constraints on their comparative belief ordering.

A further worry remains regarding the conjunction fallacy example. With just an approximation explanation we can explain how non-ideal agents can have degrees of belief, but there is still the problem that we cannot give cardinal information about an agent who has non-monotonic beliefs such as those described above. There is still the intuition that one agent David might think $A \cap B$ is much more likely than A while Arthur might think $A \cap B$ is only slightly more likely than A. However, given the measurement analogy and the requirement for monotonicity we cannot say that. This can be resolved however by examining the assumptions in the measurement analogy. On the measurement analogy, we point to being able to add lengths of objects together. Similarly then, we need to point towards an 'addition' operation for our comparative beliefs.

In section 2.4.2 I noted that if our comparative belief ordering can be represented by a probability function then we can use the union of disjoint propositions as an 'addition' operation. Using this operation binds the account to Probabilism. We can however opt for a more general requirement on an addition of beliefs operation that does not. Elliott (2020a) develops two accounts of 'Ramseyan comparativism' which allow us to retain the measurement analogy explanation for cardinality without a commitment to Probabilism.[27] The Ramsey functions Elliott defines give the ratio principle as a special case, so by moving to this weaker concatenation operation we can still get the probabilistic account.[28]

2.5.3 Interpersonal Comparability

Another desideratum for a theory of degrees of belief is whether we can express the fact that an agent Ann is more confident that *A* than an agent Bob is, i.e. can we express interpersonal facts about belief. One worry we might have about this in relation to comparative belief orderings is that as well as giving probabilistic representations of belief we can also give non-probabilistic representations. If different agents beliefs are numerically represented on different scales then it looks like we cannot make interpersonal comparisons between their beliefs.

[27] Elliott (2020a) and (Weatherson, 2016b, p.15) take Ramsey (1929) to be pointing towards comparativism that is weaker than probabilistic comparativism.

[28] Once we consider non-ideal agent's we might also be interested in agent's who's comparative beliefs cannot be represented with a probability function at all. That is, agent's who's comparative beliefs fail to meet many of the conditions for probabilistic representation that I outlined in the previous sections. In this section I have pointed towards ways we can approximate non-ideal agent's attitudes with probability functions. Another possibility is to consider the (weaker) conditions comparative beliefs must meet to be represented by a real-valued function (rather than a probability function). Konek (2019) gives a detailed overview of the (weaker) conditions. If non-ideal agent's have comparative beliefs that fail to meet several conditions for probabilistic credences it may be better to consider these weaker conditions. See (Konek, 2019, §5.3).

What is 'real'

As noted above it is possible that there are multiple probability representations of the same comparative ordering and as per the example of Andi's beliefs at times t_1 and t_2 even when it seems like there might be real differences in Andi's confidence we cannot say this is real since it is the underlying comparative belief ordering that is real. This suggests a broader question of what do we want to take as 'real' about degrees of belief. Zynda talks about what it means to really have certain degrees of belief.

> An important descriptive question arises whether being representable as having certain degrees of belief, as described by a probability function p, is sufficient for really having those degrees of belief. (Zynda, 2000, p.49)

Clearly, if there are multiple probability representations of the same comparative ordering then it is not the case that being described by a probability function is sufficient for really having certain degrees of belief. Furthermore, it's not just that there may be multiple probability representations of the same comparative ordering. Since the comparativist holds that probabilities are not what's real but rather probabilities represent the comparative beliefs it follows that if there is a probability representation then there can also be non-probabilistic numerical representations. This looks potentially problematic given certain desiderata. Can we retain Probabilism (that a rational agent's degrees of belief ought to in some sense adhere to the probability axioms) and can we retain interpersonal comparisons?

Zynda outlines different positions one might take towards degrees of belief: eliminativism, anti-realism, weak realism and strong realism.

- Strong realism: degrees of belief (and preferences and utilities) are independently existing interacting mental states.
- Weak realism: we can attribute degrees of belief to an agent and use them to "describe aspects of their psychological states, but ...degrees of belief ...should not be thought of as independently existing"(p.55). Rather, they are derivable from comparative beliefs.
- Anti-realism: We can use degrees of belief in a formal theory but we should not take them to be psychologically real.
- Eliminativism: degrees of belief are not real and we should not make use of them in our theories.

Some of these positions do not look compatible with the comparativist approach. In particular, strong realism and eliminativism. The comparativist takes comparative beliefs to be real and degrees of belief to be representations of the comparative belief structure - the degrees of belief are not thought of as independently existing and psychologically real on this picture. From the desiderata I outlined, eliminativism seems inadequate since an account ought to explain the functional role degrees of belief play in formal theories. Weak realism and anti-realism both seem compatible with the idea of there being multiple numerical representations of an agent's comparative ordering. The question still stands, what can we take as 'real' across multiple numerical representations? Stefánsson argues that if all numerical representations of a comparative belief ordering share a feature then we should take that as real.

> We should, I contend, accept as real any feature that is shared by all models of a real phenomenon. (Stefánsson, 2017, p.8)

Following this, we can say that in relation to questions about whether comparativism can capture facts about cardinality or interpersonal comparisons we need to show that non-probabilistic representations can account for the same essential features.

In the following section I will outline an example of a non-probabilistic numerical representation of beliefs.

Non-probabilistic numerical representations

An example of a non-probabilistic measure is given by Zynda (2000) who considers a function he calls a believability function and shows how both probability and believability functions can be used to represent the same comparative information. With a believability function, the agent's beliefs range from 1 to 10 (rather than 0 to 1 with a probability function). The axioms of the believability are as follows:

Definition 2.3. Let $B : \mathcal{F} \to \mathbb{R}$. B is a believability if and only if for any event $E_i \in \mathcal{F}$

B1. **Minimality**: $B(E) \geq 1$

B2. **Maximality**: $B(\Omega) = 10$

B3. **Subadditivity**: $B(E_1 \cup \cdots \cup E_n) = [B(E_1) + \cdots + B(E_n)] - (n - 1)$ if for every $i,j \in \{1, \ldots, n\}$ such that $i \neq j$, $E_i \cap E_j = \emptyset$.

The believability function is notably different from a probability function in that it is not additive, it is subadditive.[29] We can imagine two agents that have the same comparative beliefs, but one represents their comparative beliefs with a probability function and the other with a believability function. Zynda considers the case of having beliefs about the rolls of a fair die. There are 6 possible outcomes which are all equally likely and mutually exclusive so each outcome has believability ranking B(exactly one of $1, 2, 3, 4, 5$ or 6 comes up) $= 2.5$, the believability ranking that an even number comes up is B(even number) $= 5.5$. With a probability ranking we would have that P(an even number comes up) = P(a $2, 4$ or 6 comes up) but the believability ranking that a 2 a 4 or a 6 comes up is $B(2, 4$ or 6 comes up) = 7.5. The agents will have the same comparative belief ordering, so it looks like we cannot say degrees of belief obey Probabilism (since the believability function is just as good a representation of the comparative belief orderings as the probability function). This looks problematic if one holds the view that degrees of belief *are* probabilities (and obey Probabilism).

One option here might be to argue that although comparative beliefs can be represented by believability functions, this is an inferior representation to a probability function. If there is a reason to give

[29]Subadditivty in general, also referred to as the triangle inequality, is the property of a function that the function applied to the sum of two elements of the domain is at most the sum of the function applied to each of those elements. Euclidean distance is an example of a function that is subadditive. Square roots are another common example of a subadditive function, for example $\sqrt{4+4} < \sqrt{4} + \sqrt{4}$. In the case the believability function the subadditive axiom tells us that an agent's believability in the union of events A and B is less than their believaility in A summed with their believability in B.

probability functions a privileged position in representing an agent's comparative beliefs then we can introduce this as a restriction on a rational agent. We could look to various theoretical virtues of one theory over another such as simplicity, elegance, strength etc. Zynda considers these but ultimately argues that at most these theoretical virtues can make a probability ranking more likely than a believability ranking rather than definitively showing that an agent's degrees of beliefs should only be represented by a probability ranking.

Rather than seeing the non-probabilistic representations as a problem they instead look like a feature of comparativist accounts (and non-comparativist accounts such as preference-based accounts). We should treat it as a feature of the account that comparative beliefs do not restrict us to one type of numerical representation. Relative to a probability ranking, we can still retain all of the features of belief that we associate with the degree of belief representation of belief. Indeed, this is the approach Stefánsson (2017) takes saying that:

> … it might be true that Karl believes E to degree 0.5 relative to a 0–1 scale, and also true that he believes E to degree 5.5 relative to a 1–10 scale. But, I contend, it is neither true nor false that he believes E to degree x simpliciter. So this, I contend, is the sense in which degrees of belief are not psychologically real—but also the sense in which they partly correspond to something real. The difference between a probability and a believability representation of Karl's epistemic state does not correspond to anything psychologically real; but both may correspond to something psycho-

logically real given their respective scales (p.8).

The idea of using multiple scales to track the same phenomena is also one we are familiar with from how we measure temperature. We use both the Celsius and Fahrenheit scales to track the same underlying feature: kinetic energy. Similarly, with degrees of belief, we can use either probabilities or believabilities to represent an agent's beliefs and that's fine so long as they are both scales that correspond to comparative beliefs (which are what's real).

Interpersonal comparisons

This is related to the non-uniqueness problem discussed in the previous section. If an agent Ann has their comparative beliefs represented by a probability ranking and agent Bob has their comparative beliefs represented by a believability ranking then any attempts at comparing their beliefs would be meaningless. Meacham and Weisberg (2011) note that without interpersonal beliefs we cannot explain why one agent assents to a proposition while another does not (see p.659). We want to be able to say that Ann is not only more confident in a proposition A than Bob is but also to make stronger claims such as agent Ann is much more confident in a proposition A than Bob is.

Stefansson's approach then is one that we can apply to the problem presented above regarding the possibility of agents beliefs being represented by different numerical scales: we can stipulate that in order to make a meaningful interpersonal comparison we require that both agents are working on the same scale. We do this by requiring the extreme points of agents' beliefs to agree i.e. we need for Ann and Bob that they agree on the degree of belief given to a tautology and the degree of belief given to a contradiction. If both agents have degree of belief 1 about tautologies and degree of belief 0 about contradictions then their comparative beliefs will fall on the same type of model (a probability function) and we can make interpersonal comparisons of their other beliefs since we can say that Ann's belief in A is closer to her belief in a tautology than Bob's belief in A thus putting the beliefs on the same type of scale.

There are some obvious worries with this type of response to the problem of interpersonal belief. A lot is riding on the idea that we can compare the maximum and minimum points of Ann and Bob's beliefs. Elliott (2020c) argues that we cannot simply assume as Stefánsson does that if a proposition A is at the top of Ann's comparative belief ranking ($\succsim_a$) and a proposition B is at the top of an agent Bob's comparative belief ranking ($\succsim_b$) then Ann and Bob's confidence in A is equal (and similarly for a proposition B at the bottom of both Ann and Bob's comparative belief rankings). Elliott (2020c) refers to this as MIN-MAX EQUALITY. If we assume this then we can explain what it means for Ann to be more confident in a proposition R than Bob is in R. Let C_a be a numerical ranking of Ann's comparative beliefs and let C_b be a numerical ranking of Bob's comparative beliefs. Then under the assumption that they are on the same scale (they have the same numerical endpoints to their scales) we can explain what it means for Ann to be more confident than Bob is in a proposition R by looking at the comparative distance between $C_a(\top)$ and $C_a(R)$ and $C_b(\top)$ and $C_b(R)$.

The problem for the comparativist then is that we need to justify the assumptions being made. As Elliott notes:

> if MIN-MAX EQUALITY is true, then interpersonal confidence comparisons might be meaningful under certain conditions. …What comparativists need is a justification for MIN-MAX EQUALITY. (Elliott, 2020c, p.7)

It looks like we need further justification that C_a and C_b 'belong to the same model' and that both agents give the same values to the maximum and minimum. We can simply stipulate that in order to make interpersonal comparisons we need to be on the same scale (i.e. that both agents comparative beliefs are represented by probabilities or both agents comparative beliefs are represented by believabilities). We need them to be on the same types of functions, but we also need justification that the endpoints they assign on these are the same. What we need is something that justifies thinking people give the same value to at least two points on the scale. One thought is to simply stipulate that the endpoints are the same between agents. This lacks justification however. Just because we stipulate that the endpoints are the same and carry out the process given by Stefánsson does not give any reason to think that we are making meaningful comparisons. Elliott (2020c) for example gives the analogy of stipulating that the endpoints for volume and mass scales are the same. Once we have made this stipulation we can make comparisons between volume and mass. The stipulation alone does not mean that the comparisons we are making are meaningful.

The endpoints of an agent's comparative beliefs seem importantly different from Elliott's mass/volume example, however. Importantly the endpoints of comparative beliefs seem to play psychologically similar roles for agents. Although this is a step in the right direction, we can again see that playing psychologically similar roles is not enough justification by itself since the same could be said of preferences and we can't say we can make interpersonal comparisons of preferences just due to psychological similarity.[30] One way we might try to resolve this is by looking at the functional role of belief in a theory to show that the endpoints of agent's comparative belief orderings play the same role. We might for example appeal to broader features of partial belief such as agents' choice behaviour in similar situations which can do the job of providing a scale-independent relation between the beliefs of Ann and Bob.

Elliott (2020c) argues that this is not a fruitful strategy for the comparativist. He argues that appealing to other features of partial belief undermines the idea that comparative beliefs are fundamental and that we can give an explanation about what degrees of belief are from comparative beliefs. If we look at the functional role of belief in decision making then we could put the numerical degree of belief into our chosen model of decision making and from this learn an agent's utilities. The worry is that if we give an argument that identifies an agent's maximum confidence and minimum confidence as being the same in virtue of them playing the same functional role with respect to utilities then there's no reason to stop there and we

[30] See (Elliott, 2020c, §3).

could remove comparative degrees of belief from the picture entirely and just talk about the functional role of an agent's disposition to behave in a certain way. While I agree that appealing to the functional role of belief in decision making undermines the comparativist position the overall strategy of appealing to the functional role of beliefs looks like a promising strategy.

There are two strategies that could be used by the comparativist to justify that there are at least two points that different agents agree on. I will consider both in turn. One way is to appeal to the similar psychological role of the maximum and minimum beliefs of an agent's comparative beliefs in the context of the role they play in an agent's response to new evidence. Another is to say that if we want to justify interpersonal comparative beliefs we will need to appeal to something beyond comparativism. In this case, we accept that comparativism cannot explain interpersonal comparisons of belief. Although this was one of the desiderata I outlined in section 2.3.2 we might accept this sort of move on the grounds that other accounts of degrees of belief also face a similar problem with interpersonal degrees of belief. In particular, it seems accounts that appeal to utilities as part of an explanation of degrees of belief face problems. List (2003) for example talks about the orthodox view in economics being that we cannot make empirically meaningful interpersonal comparisons of utility. A clear articulation of the problem of comparing utilities can be found in Robbins (1932) where he argues "[i]ntrospection does not enable A to measure what is going on in B's mind, nor B to measure what is going on in A's. There is no way of comparing the satisfactions of two different people" (p.123).

We can appeal to something beyond comparativism by supplementing comparativism - that is we say rather than partial beliefs being determined solely by an agent's comparative beliefs we explain partial beliefs in terms of an agent's comparative beliefs and something else (provided the addition does not undermine comparativism itself). We could supplement comparativism to give an explanation for interpersonal comparisons by looking at the content of beliefs. We can then appeal to the idea that rational agents have the same level of confidence in propositions about chancy events. Unlike with the suggestion of the appeal to the functional role of degrees of belief in decision theory I do not think this undermines the main position of comparativism. We still explain degrees of belief in terms of an agent's comparative beliefs (and we cannot reduce the explanation in terms of something else like preferences). With this approach, we do have to admit a limitation of comparativism by itself and its ability to account for interpersonal beliefs. Supplementing comparativism in this way is a means to justify why we can say two agents agree on points in their comparative belief orderings rather than being used to explain degrees of belief.

We do not have the extremal points of comparative beliefs in isolation from the network of our other beliefs. We can consider Ann is just as confident in a proposition A as its negation which gives a midpoint in her confidence ranking (assuming she is rational). We can also imagine that Ann and Bob both have beliefs about chance-based events such as the likelihood of a fair coin landing heads. Given these chance events Ann can talk about her belief in A in relation to her belief in a fair coin landing heads (she thinks they are equally likely) and Bob can talk about his belief in relation to his belief that a fair die lands on 1 (he thinks they are equally likely). So long as agents also have beliefs about shared propositions where they have rational beliefs towards those propositions we can always just situate their other beliefs in relation to these. Thus allowing us to make sense of interpersonal belief comparisons. A more fine-grained example might be that Ann thinks getting good coffee from her local independent coffee shop is just as likely as a black ball being drawn at random from an urn that contains 100 balls, 90 of which are black. Bob on the other hand thinks the chance of him getting good coffee from that coffee shop is as likely as a fair coin landing heads. Indeed, a similar idea of how we make sense of a probabilistic representation of degrees of belief can be found in Ramsey (1929)

> … the notion of a belief of degree $2/3$ is useless to an outside observer, except when it is used by the thinker himself who says “Well, I believe it to an extent $2/3$”, i.e. (this at least is the most natural interpretation) ‘I have the same degree of belief in it as in $p \vee q$ when I think p, q, r equally likely and know that exactly one of them is true’. (Ramsey, 1929, p.256)

A worry with this approach is that we need to assume that agents are rational/ it seems like we can only explain comparisons between rational agents. Moreover, it requires the assumption that all rational agents will have the same credence towards chancy events. For example, agents might have different attitudes due to misleading evidence. This approach does not require that all agents have the same attitude to all chance events however, merely that if two agent's wish to compare their beliefs, then they must have the same attitude towards at least two chance propositions. Just as the measurement analogy requires that an agent have enough beliefs for the analogy to work, we can also put this constraint on interpersonal comparisons. We can then consider only propositions where an agent has admissible evidence and they can apply a chance-credence principle.

While appealing to the content of an agent's beliefs looks promising, there is another approach we might take that does not supplement comparativism. We can look at how agents respond to evidence. We do not need to appeal to utilities or the role of degrees of belief in decision theory to show the functional role of beliefs. We instead can see how agents update their existing comparative beliefs. When an agent updates her beliefs by conditionalization we know that conditionalization maintains certainties. If an agent is certain of a proposition (such as a tautology) at a time t_1 then if she gets new evidence and updates on this by conditionalization she will remain certain in that proposition at time t_2. Taking this idea we can apply it to the comparative belief ordering and can explain what it means for a belief to be at the top of a comparative belief ordering. If beliefs at the top and bottom of Ann's comparative belief order do not change with new evidence and similarly for Bob's then we have a justification that what it is for Ann and Bob to have 100% and 0% confidence in terms of the role it plays in updating on new evidence.

This shows how we can show that agents have the same degrees of belief by looking at the functional role they play in updating on new evidence. As well as looking at how an agent updates her beliefs by conditionalization there are also accounts that suggest we look at the stability of belief. Stability of belief is a way of accommodating the distinction between ways evidence can influence belief. Joyce (2005b) notes that there is a distinction between the weight and balance of evidence we might get. The distinction can be seen if we consider the example of two coins. One we know is a fair coin and with the other we don't know if it is fair or biased. With the fair coin, you see 1000 tosses of the coin and it comes up heads about half the time. Here there is weight of evidence that the coin will land heads 50% of the time. With the coin where you're uncertain whether it's biased and you've not seen any tosses of the coin. Here on the balance of evidence it seems like you should be no more confident in the coin landing heads or tails, so the balance of evidence suggests you take it that the chance of the coin landing heads to be 50%. So, we can give a story that explains how an agent might have the same degree of belief in both these cases but it seems like there is something importantly different in the evidence in the two cases. One way that it has been suggested that we can show this difference is by looking at the stability (Leitgeb (2014)) or resiliency (Skyrms (1977)). Leitgeb uses stability to characterise the connection between belief and degrees of belief. Namely that belief corresponds to resiliency of high probability.

This suggests a problem for my account since it seems that an agent could give less than 100% confidence to tautologies and 0% to contradictions but her beliefs play the same functional role as an agent that does give 100% confidence to tautologies and 0% to contradictions. We can imagine an agent Adina who has comparative beliefs but she holds that one should never be fully certain of anything. As such she has some doubt towards tautologies and some slight doubt towards contradictions. Adina's beliefs in tautologies and her beliefs in contradictions are resilient. In this case, even though she does not assign credence 1 (or credence 0) to them respectively, she resiliently has beliefs that stay at the top and bottom of her comparative belief ranking.

While resilient beliefs seem to play a similar functional role if an agent has resilient belief rather than credence 1 it is possible that, given enough evidence she will change her beliefs. So, there is a distinction in the functional role. The example of Adina seems problematic because it seems the circumstances around why she does not give 100% confidence to tautologies and 0% to contradictions make her beliefs more than just resilient but unchangeable. This seems wrong, however, if she is not fully certain of anything this includes her commitment to never being fully certain. There is nothing that removes the possibility of her changing her mind and giving full confidence to tautologies.

2.6 Conclusion

In this chapter I have shown that comparativism is a viable approach to explaining what degrees of belief are. In section 2.3 I considered what conditions must be met for an account of degrees of belief to be adequate. The rest of the chapter shows that the desiderata I give in section 2.3.2 can be met by a comparativist account. In section 2.4.1 I outlined how we can give a probabilistic representation of comparative beliefs. The measurement analogy shows that we can make sense of the idea of extracting cardinal information from an agent's comparative belief ordering.

I then considered some problems that have been raised for comparativist accounts. The problem of non-unique probability measures, the possibility of comparative beliefs not meeting the conditions for probabilistic representation and the problem of whether we can give interpersonal comparisons of belief. I have shown that comparativism can respond to these problems and therefore meet the desiderata for an account of degrees of belief.

Showing that comparativism is a viable account is the first stage in justifying an intersectionist interpretation of imprecise credences. This interpretation is needed to make sense of the suspended judgment attitude I have argued for in chapter 1. As outlined in section 2.2 many of the accounts of the intersectionist interpretation assume or draw on comparativism, so showing comparativism is a viable view to hold supports the viability of the intersectionist interpretation. In the next chapter, I will focus on the second part of justifying an intersectionist interpretation of imprecise credences arguing that once we have accepted a comparativist account of degrees of belief, intersectionism is the most natural way of representing that an agent has an incomplete comparative belief ordering (rather than other interpretations of imprecise probability).

Chapter 3
Intersectionism and Suspended Judgment

3.1 Introduction

In the previous chapter I outlined the comparativist intersectionist interpretation of imprecise credences. This interpretation uses comparativism as part of its interpretation of the credal set and proposes that we use the credal set to represent an agent's determinate attitudes. Comparative intersectionism represents an agent's beliefs in terms of their comparative beliefs and also allows us to represent that an agent lacks certain comparative beliefs. This interpretation, therefore, allows us to represent that an agent has incomplete beliefs (and incomplete comparative belief orderings).

As noted in the previous chapter, to defend a comparativist intersectionist interpretation, there are two aspects of the interpretation that need defending in turn. The comparativist aspect of the interpretation where we look at an agent's comparative beliefs, and the intersectionist interpretation of those comparative beliefs. In the previous chapter I have shown that comparativism is a viable account of degrees of belief by showing that, given an agent's comparative degrees of belief we can give a probabilistic representation of their beliefs.

There are a range of different interpretations of imprecise credences and in this chapter I defend the intersectionist interpretation as a way of representing an agent's comparative beliefs over another interpretation - the supervaluationist interpretation. Supervaluationism is another prominent interpretation of imprecise credences. In the supervaluationist interpretation the credal set is interpreted as representing that an agent has vague or indeterminate attitudes. These two different interpretations are both compatible with comparativism. However, they represent different things about an agent's attitudes.

With a comparativist picture, a desirable feature of a formal representation is to be able to account for agents sometimes lacking an opinion between two propositions. Indeed, accepting anything like a comparativist (or preference-based) account of degrees of belief immediately suggests rejecting the completeness constraint: the requirement that for any pair of propositions A and B an agent is either more confident in A than B, more confident in B than A or equally confident in A and B. Completeness is an unrealistic constraint, even for ideal agents, since the agent may have insufficient evidence. It should be possible for a model of belief to account for an agent having gaps in their comparative rankings of propositions. Having an incomplete comparative ordering seems like a more realistic picture of an agent's degrees of belief. There is, however, more than one way this incompleteness could come about. One thought is that agents have incomplete comparative belief orderings because they genuinely lack certain comparative beliefs - in this case, an agent's determinate beliefs are incomplete and we can represent this with the intersectionist interpretation. Another way an agent might have incomplete comparative beliefs is due to vagueness or indeterminacy in their beliefs. As such, comparativism is also compatible with a supervaluationist interpretation. Accepting a comparativist account is therefore not enough to justify the intersectionist interpretation of imprecise credences.

As I argue in Chapter 1, there are a number of different motivations for adopting imprecise credences. In section 3.2 I will review the incompleteness motivation and why we might want to represent that someone lacks an opinion. In section 3.3 I will outline the supervaluationist interpretation and how, from the picture of comparativism I have argued for in the previous chapter, this view is compatible. After outlining a supervaluationist interpretation I will give a more detailed explanation of comparativist intersectionism and how it differs from supervaluationism in section 3.4. I will then argue that comparativism more naturally fits with this view. After I have defended the use of the comparativist intersectionist interpretation of imprecise credences I will revisit the suspended judgment attitude I discussed in Chapter 1. In section 3.5 I will show how we can understand suspended judgment within the intersectionist interpretation.

3.2 Imprecise Credences

As noted above, one reason why imprecise credences look desirable is that they enable us to account for agents having incomplete comparative belief orderings. Precise credence accounts also do not provide a clear distinction between indifference and undecidedness. As Kaplan (1996) notes:

> Both when you are indifferent between *A* and *B* and when you are undecided between *A* and *B* you can be said not to prefer either state of affairs to the other. Nonetheless, indif-

> ference and indecision are distinct. When you are indifferent between *A* and *B*, your failure to prefer one to the other is born of a determination that they are equally preferable. When you are undecided, your failure to prefer one to the other is born of no such determination. (Kaplan, 1996, p.5)

Many others have also argued for the need for imprecise credences even for ideal agents in cases of lack of evidence. We might, for example, have evidence both for and against the chance of rain "the barometer is high, but the clouds are black" (Keynes, 1921)[p.31f]. In this case, we might think there is no admissible way to weigh different types of evidence up against each other. Rabinowicz argues in such a case it "might be rational of us to remain in the state of a credence gap" (p.3916). He uses the term credence gap to mean that two propositions *A* and *B* are incomparable, and we cannot say that *A* is given more credence than *B* or vice versa and we also cannot say that they are given equal credence. We might also refer to this position as saying that the agent is undecided between *A* and *B*.

Kaplan (2010) also argues for credence gaps saying that:

> (i) you should rule all and only the assignments the evidence warrants your regarding as too high or too low, and you should remain in suspense between those degree of confidence assignments that are left, where
>
> (ii) it is a condition on a degree of confidence assignment's being among those that are left that it satisfy the axioms of the probability calculus. (Kaplan, 2010, p.45)

We cannot accommodate credence gaps in a precise credal framework since precise credences require completeness. Therefore, we ought to look towards imprecise credences if we want to give a more psychologically realistic model of numerical degrees of belief.

There is then the question about what features of the credal set represents an agent's belief state. Indeed, we can see there are features of the model which do not represent an agent's attitude; each of the credence functions in the set are precise but a key feature of the credal set is that the agent does not have a precise credence. Moreover, with the different interpretations I will outline in sections 3.3 and 3.4, we can see that from the same set of credence functions we can read off different information about the agent's beliefs depending on the interpretation. While there are a number of accounts of imprecise credences they have some features in common; they agree that there are circumstances where an agent doesn't have a precise credence. We can also see that precise credences fall out as a special case of a credal set interpretation. When a set of credence functions is comprised of just one credence function we are in the precise case.

Two popular interpretations of imprecise credences are the intersectionist interpretation and the supervaluationist interpretation. The intersectionist interpretation takes the whole credal set as representing the agent's beliefs. In contrast, the supervaluationist interpretation is an interpretation of how to model an agent's credences in cases of indeterminacy.

3.3 Supervaluationism

The supervaluationist interpretation of imprecise credences is one that holds that sometimes it is indeterminate what our beliefs are and that we should interpret indeterminacy in a broadly supervaluationist way. That is, in a way akin to a supervaluationist approach to vagueness.[1] 'Broadly supervaluationist' is a vague characterisation of the position, and as Elliott (ms) notes, this 'broadly supervaluationist' treatment results in different varieties of credal supervaluationism (see van Fraassen (1990), Hájek (2003), Rinard (2017)). For the purposes of this account, it will suffice to give a general account of credal supervaluationism that gives a description of the common features of different credal supervaluationist interpretations.

Credal supervaluationism interpretation applies more naturally to cases of indeterminacy. It is this notion of indeterminacy that is central to the distinction between the intersectionist interpretation of imprecise credences and the supervaluationist interpretation of imprecise credences.

In a credal supervaluationist interpretation we consider a credal set C comprising of credence functions: $C = \{Cr_1, Cr_2, Cr_3, \dots, Cr_i\}$. Each of the credence functions in the set are permissible precisifications of the agent's belief state. It is indeterminate which of these credence functions represents the agent's beliefs. We can say that if for a proposition A, $Cr_i(A) = x$, $\forall Cr_i \in C$ then the agent's credence in A is determi-

[1] See Elliott (ms).

nately *x*. i.e. if a proposition is assigned the same credence by every credal function in the credal set, then the agent has a determinate attitude towards this proposition.

There are strong parallels between the supervaluationist interpretation of credal sets and the supervaluationist approach to vagueness. The supervaluationist approach to vagueness places indeterminacy as a linguistic phenomenon: there are multiple possible precisifications/interpretations of the language. In this framework, borderline statements lack a truth value. For example, consider the statement: 'Bob is tall'. Let us assume that Bob is a borderline case of tall. There are some admissible ways of making the term 'tall' precise and on some of these it comes out as true that Bob is tall and on others, it comes out as false that Bob is tall. Our use of language governs what is an admissible sharpening or precisification of 'tall'. Vague terms are characterised by the fact that there is not a consensus about the admissible precisification of the term. In this case 'Bob is tall' is neither true nor false, so some features of classical logic such as bivalence do not hold. However, we can still retain some nice features of classical logic on the supervaluationist approach since it is still true that 'Either Bob is tall or Bob is not tall' on all admissible precisifications of the term tall, i.e. the law of the excluded middle holds for the object language. When a sentence is true on all admissible precisifications we say that it is super-true (or determinately true) and when it is false on all admissible precisifications we say it is super-false (or determinately false).

In the doxastic case, we need to weigh up the demands for accuracy and specificity. When we consider vagueness or indeterminacy it seems wrong to have a precise attitude towards borderline cases. Consider, for example, a paint chip, patchy, that is a borderline case of red/orange. Is patchy red? What attitude should an agent adopt towards this? One motivation for credal supervaluationism is that a credal supervaluationist approach can represent that an agent has indeterminacy in their attitude. An agent's attitudes are represented by an imprecise credal set where each of credence functions is an admissible precisifications of the agent's credence. When these functions disagree about the precise credence a proposition is assigned we take the credal set as representing that an agent has an indeterminate attitude towards it. In this example, we can represent that an agent has an indeterminate attitude about whether patchy is red if they have a credal set and some credence functions in the set give high credence to patchy is red and other credence functions give low credence to patchy is red.

That is not to say that the supervaluationist interpretation of the credal set is the only way we might represent an agent's attitudes towards vague propositions Mahtani (2019b) for example offers an alternative account based on the idea that the term 'credence' is vague. In chapter 5 I discuss the possible range of attitudes that one could take towards indeterminacy in more detail. While it may be natural to talk about credal supervaluationism as representing the attitude one takes towards borderline cases, we can also talk about indeterminacy more generally. We can distinguish between an agent having indeterminate attitudes and an agent having attitudes about borderline propositions.[2] I am interested in the supervaluationist interpretation of the imprecise credal set as a way to represent that it is indeterminate what an agent's attitude is (rather than only considering an agent's attitude towards vague propositions). One could have indeterminate attitudes towards sentences where there is no vagueness.

We can also take other information from the supervaluationist interpretation of the credal set. For example, consider the following three toy examples of credal sets for agents A, B and C. Where *C* is used for a credal set and *Cr* for a precise credence function.

Agent A:

$$C_A = \{Cr_{a1}, Cr_{a2}, Cr_{a3}\}$$

where

$$Cr_{a1}(q) = 0.6, Cr_{a2}(q) = 0.6, Cr_{a3}(q) = 0.6$$

Agent B:

$$C_B = \{Cr_{b1}, Cr_{b2}, Cr_{b3}\}$$

where

$$Cr_{b1}(q) = 0.6, Cr_{b2}(q) = 0.7, Cr_{b3}(q) = 0.75$$

and agent C:

$$C_C = \{Cr_{c1}, Cr_{c2}\}$$

where

$$Cr_{c1}(q) = 1, Cr_{c2}(q) = 0$$

[2]On some views of the cognitive role of indeterminacy these will end up being the same. I talk about the cognitive role of indeterminacy in more detail in Chapter 5.

For agent A all credence functions in her credal set assign credence 0.6 to a proposition q so we can say that she has credence 0.6 in q. For Agent B not all credence functions in her credal set agree so this represents that she has indeterminate beliefs towards q. We can still interpret the set as giving some representation such as agent B has credence higher than 0.6 in q. For agent C there is no consensus between her credence functions- she has maximally indeterminate beliefs.

As with traditional supervaluationism, there is a problem analogous to high-order vagueness; when we give an interval or set of credences rather than giving one credence, we represent the interval by two end-points. An obvious issue is that there often appears to be insufficient evidence to justify picking any particular end-points. I discuss higher-order vagueness and higher-order indeterminacy for credal supervaluationism in more detail in the next chapter.

The intersectionist interpretation, on the other hand, can be thought of as representing an agent's belief state in cases where there is no indeterminacy.

3.4 Intersectionism

In this section, I give a more detailed look at how we can give an intersectionist account and in particular a comparativist intersectionist account. With an intersectionist interpretation of an imprecise credal set, we do not interpret the credal set as containing credence functions that are permissible precisifications of the agent's credence. Rather the credal set as a whole represents the agent's determinate attitude.

Indeed, it is determinate that none of the credence functions in the credal set represents the agent's belief state.

There are a variety of intersectionist interpretations of credal sets given by Kaplan (2010), Nehring (2009) and Alon and Lehrer (2014). The intersectionist interpretation naturally follows as a way of representing incomplete comparative belief orderings and all of these authors have assumed or argued for comparativism as the basis of their accounts:

> we propose to model probabilistic beliefs as a comparative likelihood relation $\trianglerighteq$ over events, with "$A \trianglerighteq B$" denoting the judgment "A is at least as likely as B." (Nehring, 2009, p.1055)

> statements of the type 'event A is more likely than event B' seem fundamental [...] a likelihood relation [...] may be incomplete. (Alon and Lehrer, 2014, p.477)

Kaplan (2010) also spells out his interpretation in terms of an agent's comparative confidence in propositions A and B.[3] I will outline Kaplan's interpretation in more detail below.

There are some features of the model that we take to represent the agent's belief. If all the credence functions in the set agree that two propositions A and B have the same credence, then we interpret this as representing that the agent thinks A and B are equally likely. The following examples from Elliott show more precisely what can be read from the credal set:

[3]Notation altered for consistency.

$Cr_i(A) = Cr_i(B) \;\; \forall Cr_i \in C$ iff '$A \sim B$' is true.

Where $A \sim B$ iff the agent believes A and B are equally likely.

However, when we consider the conditions for saying that $Cr_i(A) \neq Cr_i(B)$. We cannot say:

$Cr_i(A) \neq Cr_i(B) \;\; \forall Cr_i \in C$ iff '$A \sim B$' is false.

Since the agent might have a credal set such that for some $Cr_i \in C$ $Cr_i(A) = Cr_i(B)$ but there is some Cr_j such that $Cr_j(A) \neq Cr_j(B)$. This leads to the condition that $Cr_i(A) = Cr_i(B)$ iff the agent is equally confident in A and B is both true and false.[4] (Elliott, ms, p.12)

This shows the need for a clear interpretation of what aspects of the credal set we take to represent an agent's beliefs. There are a number of ways we can interpret the credal set. Kaplan (2010) gives a version of intersectionism he calls Modest Probabilism. He characterises the confidence assignments for a credal set C and for every pair of propositions A and B.[5]

(i) you are equally confident in A as you are in B if and only if every member of C assigns A the same value as it assigns B;

(ii) you are more confident in A than you are in B if and only if every member of C assigns A at least as great a value as it assigns B, and at least one member of C assigns A a greater value than it assigns B; and

[4] Notation altered for consistency.

[5] Notation altered for consistency.

(iii) otherwise you are undecided as to the relative credibility of A and B.

From Kaplan's characterisation we can see that there are some general things we can say about an agent's credences based on the credence functions in the credal set. If an agent α has a credal set $C_\alpha = (Cr_1, Cr_2, Cr_3, \dots, Cr_n)$ then

- if $\forall Cr_i \in C_\alpha,\ Cr_i(A) = k$ then we can say that α has degree of belief k in A (i.e. $C_\alpha(A) = k$).
- if $\forall Cr_i \in C_\alpha,\ Cr_i(A) > k$ then we can say that α has degree of belief greater than k in A (i.e. $C_\alpha(A) > k$).
- Kaplan's interpretation also means we can say that $C_\alpha(A) > k$ if $\forall Cr_i \in C_\alpha,\ Cr_i(A) \geq k$ and $\exists\, Cr_j$ such that $Cr_j(A) > k$.

and so on.

The following examples show in more detail what an agent's comparative beliefs represent on this interpretation.

> Example 1: If we consider the credal set of agent α, $C_\alpha = \{Cr_i, Cr_j\}$ where $Cr_i(A) \geq Cr_i(B)$ and $Cr_j(A) \geq Cr_j(B)$. Then by the above interpretation $A \gtrsim B, A \succ B$ and $A \nsim B$.
>
> Example 2: If we consider the credal set $C_\alpha = \{Cr_k, Cr_l\}$ where $Cr_k(A) > Cr_k(B)$ and $Cr_l(A) < Cr_l(B)$. Then by the above interpretation $A \nsim B, A \nsucc B$ and $B \nsucc A$ so α is undecided about the relative likelihood of A and B.

There are other interpretations of credal sets that give different versions of intersectionism. However, for the purposes of providing an interpretation of suspended judgment, taking Kaplan's interpretation of Modest Probabilism will suffice.

As we can see, the version of intersectionism I have presented above is given in comparative belief terms and thus it looks natural to give this account if one holds a comparativist view. It is however possible to be a comparativist and hold only a supervaluationist interpretation. Rather than interpreting the incomplete comparative ordering via the intersectionist interpretation we could instead view incomplete orderings as the product of vagueness in the comparative ordering. This vagueness can then be represented by an imprecise credal set with a supervaluationist interpretation.

One needs a reason for thinking that agents have incomplete comparative belief orderings, not just indeterminate or vague comparative beliefs. From the formalism alone we can't give an argument one way or the other. The formal structure of the set is the same for the supervaluationist and intersectionist interpretations. The interpretations are just telling us philosophical stories about what we can say is representational about the structure. We also need *both* a supervaluationist and an intersectionist interpretation of imprecise credal sets. Indeed, if the intersectionist interpretation is representational of an agent's determinate attitudes then this leaves open the question of how to represent that it may be indeterminate what an agent's attitude towards a proposition is. The supervaluationist interpretation of a credal set is a natural way of representing this. So, even though I'm arguing for an intersectionist interpretation that represents an agent's determinate incomplete attitudes, it is compatible to also argue that an agent may have indeterminate or vague attitudes which we can represent with a supervaluationist interpretation.[6]

In light of this, to argue for the intersectionist interpretation in this instance is to argue that we need imprecise credences that represent an agent's determinate attitudes as well, rather than just saying the incompleteness in a comparative ordering is vagueness. Incompleteness of belief looks like a reasonable assumption if we take certain metaphysical views about belief such as taking a view of mental representation of belief such as Fodor's language of thought (see Fodor (1975) and Rescorla (2019)) or in some inner 'language' (see Field (1978)). That is, mental representations have a central role "to believe that *p*, or hope that *p*, or intend that *p*, is to bear an appropriate relation to a mental representation whose meaning is that *p*" (Rescorla, 2019).

If we take the comparative belief ordering to be picking out some real mental representation, then gaps in the comparative ordering look natural. If there is some sense in which the comparative belief ordering has to be represented by the brain, then a complete comparative belief ordering would be far too rich (even if we take the view that some parts of the ordering are implicit). Someone can just fail to have comparative beliefs between two propositions.

[6] I discuss this in more detail in Chapter 4.

3.4.1 Problems for intersectionism

Maher (2006) argues that there are problems with accepting imprecise credences and credal sets. His argument is that for an agent to have a credal set they would have to have credences that are unwarranted by their evidence. This is exactly what we are trying to avoid when we make the move from precise credences to imprecise credences. Maher gives the following example:

> Let P be that there is intelligent life in the Andromeda galaxy. Suppose a ball is to be drawn randomly from an urn containing balls numbered from 1 to 1,000, and let Q_n be the proposition that the ball drawn has a number less than or equal to n. My degree of belief in P is higher than my degree of belief in Q_0 and less than my degree of belief in Q_{1000}. However, since my degree of belief in P is vague, there is no n such that my degree of belief in P is higher than my degree of belief in Q_n but not higher than my degree of belief in $Q_n + 1$. From these facts and [Modest Probabilism] it follows that my degrees of belief are epistemically irrational. (Maher, 2006, p.146)

While this seems to pose a problem for Kaplan's Modest Probabilism it is important to note that just because we can generate this sort of problem for imprecise credences that doesn't mean precise credences are the better option. Indeed, precise credences face problems of this form and more.

This is also an instance of a general problem for imprecise credence interpretations. We could pose this for the supervaluationist interpretation under a different description: how do we justify the apparent sharp cut-offs of upper and lower probabilities of the credal set. I will look at this problem for supervaluationism in more detail in the next chapter.

3.5 Suspended Judgment

Turning back now to suspended judgment, I will show how the intersectionist interpretation gives a natural interpretation of the suspended judgment attitude. With the account of imprecise credences that I have presented there are four possible attitudes an agent can take towards a pair of propositions. An agent can believe:

- *A* is at least as likely as *B*.
- *A* is strictly more likely than *B*.
- *A* and *B* are equally likely.
- have a credence gap for *A* and *B*.

Different interpretations of the credal set mean that the same credal set can be seen as representing different beliefs towards *A* and *B*, as shown in the examples in the previous section. For the purposes of this discussion, the particular intersectionist interpretation we choose is not important and so I will work with Kaplan's interpretation.

The credence gap attitude occurs on Kaplan's interpretation when the credence functions in the credal set disagree about the relative likelihood of *A* and *B*. For example, in the following credal set:

$C = \{Cr_i, Cr_j\}$ where $Cr_i(A) > Cr_i(B)$ and $Cr_j(A) < Cr_i(B)$.

I suggest we should understand the attitude of suspending judgment as having a credence gap. This gives an interpretation to the formal requirements on the interval that I gave on the suspended judgment in Chapter 1. In Chapter 1, I argued that suspended judgment should be understood as having imprecise credence of the form $C = [a, b]$, $a < 0.5$ and $b > 0.5$. When these conditions are met, the agent also has a credence gap towards the proposition.

When an agent has a credence gap, there is not enough consensus amongst the credence functions in the credal set. This is reflective of the type of situations that would lead to an agent having credence gaps such as lack of evidence for the agent to form an opinion about the relative likelihood of *A* or *B* or conflicting evidence. These appear to be exactly the kind of reasons why an agent would want to suspend judgment. The intersectionist interpretation of the credal set implies that a credence gap is adopted in cases where the agent does not perceive the lack of completeness in their comparative beliefs to come from vagueness but rather that they determinately think they lack an attitude with regards to this proposition. That is to say, when I suggest that an agent suspend judgment in the Probabilistic Liar what we want to represent with the suspended judgment attitude is the agent determinately lacking any precise attitude towards the likelihood of the proposition expressed by the Probabilistic Liar (rather than expressing that the agent does, in fact, have a precise attitude, but that they are unsure about which precise credence function represents it).

I will take a closer look now at how this interpretation of suspended judgment provides a solution to the Probabilistic Liar. Recall the Probabilistic Liar asks us to imagine an agent Alex who stipulates the meaning of a sentence (α) to be the following:

$$(\alpha) \quad Cr\rho(\alpha) < 0.5$$

The solution I suggested in the previous chapter was that Alex shouldn't have any sharp credence in (α) since any sharp credence in (α) leads to a conflict between the norms of Rational Introspection and Probabilism. We can also see in the intuitive reasoning about the Probabilistic Liar case that if Alex has any precise credence, and she then reflects on the evidence she gets about the Probabilistic Liar from having that precise credence, her evidence undermines her credence. This leads to a systematic instability in what credence she has.

My suggestion was that Alex instead suspends judgment about (α) and suspended judgment should take the form of having imprecise credences, in particular having a credal set that covers the entire interval which we can represent as:

$$C\rho(\alpha) = [0, 1]$$

We can see with this credal set that Alex has a credence gap for (α) and ($\neg\alpha$). Let the following describe Alex's credal set.

$$C = \{Cr_1, Cr_2, Cr_3, \ldots, Cr_n\}$$

and consider what these credence functions say about (α) (given the earlier stipulation that they range $[0, 1]$). For example, we could order Alex's credences as follows:

$$
\begin{aligned}
&Cr_1(\alpha) = 0,\\
&\vdots\\
&Cr_i(\alpha) = 0.2,\\
&\vdots\\
&Cr_j(\alpha) = 0.85,\\
&\vdots\\
&Cr_n(\alpha) = 1
\end{aligned}
$$

Recall that we've assumed that each of the credence functions is probabilistically coherent so given the above we can also say what each credence function says about $(\neg\alpha)$:

$$
\begin{array}{c||c}
Cr_1(\alpha) = 0 & Cr_1(\neg\alpha) = 1\\
\vdots & \vdots\\
Cr_i(\alpha) = 0.2 & Cr_i(\neg\alpha) = 0.8\\
\vdots & \vdots\\
Cr_j(\alpha) = 0.85 & Cr_j(\neg\alpha) = 0.15\\
\vdots & \vdots\\
Cr_n(\alpha) = 1 & Cr_n(\neg\alpha) = 0
\end{array}
$$

We can see that some of the credence functions in Alex's credal set say that (α) is more likely than $(\neg\alpha)$ and some of the credence functions say that (α) is less likely than $(\neg\alpha)$. Using Kaplan's interpretation of the credal set we can see that Alex has a credence gap for the pair of propositions (α) and $(\neg\alpha)$. When Alex is undecided or in suspense between (α) and $(\neg\alpha)$ this indecision blocks the instability between her thinking first (α) is true and then $(\neg\alpha)$ is true.

3.6 Conclusion

In this chapter I have argued that there is good reason to take the intersectionist interpretation of imprecise credences to represent an agent's determinate attitudes. In the previous chapter I argued that comparativism was a viable account of degrees of belief. Given comparativism it looks like a representation of an agent's beliefs ought to be able to accommodate that an agent can have incomplete comparative beliefs. We cannot represent that an agent has incomplete comparative beliefs with precise credences, so we need an imprecise credal model in order to accommodate this.

There are however a number of different interpretations of imprecise credences. Two major interpretations of imprecise credences are the supervaluationist interpretation and the intersectionist interpretation. There are two ways we might think incomplete comparative belief ordering come about. One way it might come about is by the agent's comparative beliefs being indeterminate. Indeterminacy in the comparative belief ordering leaves open the possibility that there is no determinate fact about whether an agent thinks A is at least as likely as B or whether they believe B is at least as likely as A or both. Another way this incompleteness might come about is if an agent lacks a comparative belief, they may simply lack an opinion about the relative likelihood of A and B. In other words, a 'gap' exists in the agent's comparative belief ordering.

I have argued that we ought to be able to represent that an agent can have gaps in their comparative belief ordering (which is also compatible with an agent having indeterminate attitudes as well). The intersectionist interpretation is a natural way of representing that an agent has gaps in their comparative belief ordering. With this interpretation, we can formally represent credence gaps.

We can represent the attitude of suspended judgment with the intersectionist interpretation of imprecise credal sets. In particular, when an agent has an imprecise credal set that fulfils the criteria for suspended judgment I gave in Chapter 1 they will have a credence gap. The intersectionist interpretation of imprecise credences therefore supports my claim that we should reject Caie's implicit assumption that an agent's determinate attitudes as represented by precise credences.

Chapter 4
Revenge Paradox

4.1 Introduction

This chapter will explore a problem generated for the account given so far- the Revenge problem. Revenge problems are a familiar problem from solutions to the Liar paradox. They offer strengthened versions of the original Liar paradox and show that even if a theory can adequately overcome the Liar paradox it is still vulnerable to paradoxical sentences. Just as we can generate Revenge problems for the Liar it seems we will also be able to generate Revenge problems for the Probabilistic Liar.

Given the parallels between the Liar paradox and the Probabilistic Liar looking to the Liar paradox and Revenge for the Liar provides a useful outline of how Revenge problems can be generated and the types of solution that can be given. In section 4.2 I outline how Revenge problems can be generated for the Liar paradox. This will provide the basis of how a Revenge problem can be generated for the Probabilistic Liar in light of the solution I have argued for (suspended judgment). In section 4.2.1 I will consider two types of response that have been given to Revenge for the Liar: denying that a Revenge paradox can be generated, and embracing Revenge and the hierarchy of metalanguages that are then needed to respond to each new Revenge problem.

After looking in detail at Revenge for the Liar and proposed solutions I will outline how a Revenge problem can be generated for the Probabilistic Liar in section 4.3. I propose that a Revenge problem can be generated if one can refer to the contents of an agent's credal set. We can then give a sentence that refers to itself and the contents of the credal set in such a way that any precise credence or the suspended judgement attitude is problematic. This suggests the need for the agent to adopt a new type of attitude towards the Revenge problem. In section 4.5 I argue that it ought to be indeterminate what attitude an agent has towards the Revenge problem. I show how this can be represented by a set of sets of credal sets which we interpret with a credal supervaluationist interpretation.

In line with the Revenge problem for the Liar, once a solution to Revenge for the Probabilistic Liar has been given we can again ask if a new strengthened Revenge problem can be generated. The solution I suggested for the Revenge problem relies on the agent not having a determinate attitude towards Revenge. In section 4.6 I show that we can generate a hierarchy of Revenge problems in terms of a determinacy operator. This can be done by considering iterations of an agent *not determinately not* having a particular attitude. That is, if we can refer to the lack of determinacy of an agent's attitude, further Revenge problems can be given. I show how we give solutions to each of these Revenge problems in terms of indeterminacy of an agent's attitude to each new Revenge problem.

In section 4.6.1 I will consider a problem that may occur based on certain assumptions about indeterminacy. Williamson (1994) shows that given certain assumptions that seem natural to make about higher-order indeterminacy we get the result that indeterminacy cuts out. This result would undermine the account I have outlined of Revenge and strengthened Revenge problems for the Probabilistic Liar. In section 4.6.2 I show how this problem can be generated and how we can avoid it if we take a weak logic of determinacy.

4.2 Revenge for the Liar

In this section, I give a detailed look at how Revenge problems are generated for the Liar. In particular, I note the underlying assumptions that must be made to generate a Revenge problem for the Liar and how the Revenge problem looks unavoidable for any solution to the Liar paradox. In section 4.2.1 I will then give a detailed look at a type of solution to Revenge for the Liar. This will inform the basis of the discussion in section 4.3 and 4.5 where I outline how we can analogously generate a Revenge problem for the Probabilistic Liar and offer a solution to the Revenge Problem I have generated.

In theories of truth we encounter the Liar paradox if we assume our language contains the resources to meaningfully self-refer and have a truth predicate. We can then consider sentences such as:

> λ: λ is false.

and we can ask whether λ is true or false. By the bivalence assumption λ must be one of true or false and by the biexculsion assumption λ can be only one of true or false. Taking each assumption in turn we see that if we assume λ is true then what it says is true, so it is false. So λ is both true and false which is a contradiction. So we now assume that λ is false, λ says it is false so what λ says is in fact the case so it is true, which again tells us λ is both true and false.

In order to avoid a contradiction we are forced to give up some platitude about truth or our logic such as bivalence, law of the excluded middle or the Tarski T-schema. Solutions to the Liar paradox have to give up or restrict at least one of these. One family of solutions to the Liar paradox gives up bivalence and considers a third category that a sentence can fall into other than true or false such as undefined (Kripke (1975)) or undetermined (Maudlin (2004)) or pathological (Cook (2007)).[1] I will focus on this type of solution to the Liar paradox since my proposed solution to the Probabilistic Liar has strong parallels.

Let us call the truth-value gap 'undefined'. While there are a number of ways of making this into a theory of truth, the important feature for this discussion regarding the Revenge problem is that we have a non-contradictory way of categorising the Liar sentence: it is not true or false, it is undefined. Using this category for the Liar sentence we can avoid the original Liar paradox. However, once we have introduced a concept into our object language that purports to deal with the Liar sentence we are open to a Revenge paradox or strengthened Liar paradox. We can use this category to formulate a strengthened Liar sentence that says:

[1] In general, we can take the idea of having a truth-value gap and use Strong Kleene semantics.

> σ: σ is either false or undefined.

As with the Liar sentence, we can give an argument to contradiction. We assume we have three semantic categories true, false and undefined which exclude each other.

> Assume σ is true, then what σ says is the case, so σ must be either false or undefined. So σ is either both true and false or true and undefined both of which are contradictions so contrary to our initial assumption σ cannot be true.
>
> Assume σ is false, which implies that σ is either false or undefined. σ says it is either false or undefined, however, so what it says is the case, so it is true. So σ is true and false which is a contradiction with the exclusion assumption.
>
> Assume σ is undefined, which implies that σ is either false or undefined. σ says it is either false or undefined, however, so what it says is the case, so it is true. So σ is true and undefined which is a contradiction with the exclusion assumption.

As before, it seems that in order to avoid contradiction we are forced to introduce a new semantic category 'other'. If we have the assumption that each category we introduce excludes the other categories then we can see that with this new semantic category we can generate another Revenge paradox in the same way as we generated the first:

> ω : ω is either false or undefined or other.

In general, each time we have a proposed solution we can generate a new problem sentence by subsuming the concept used in the solution into our object language. Any account of truth that claims to offer a solution to the Liar paradox must therefore also offer a solution to Revenge. Cook says the following about Revenge:

> Revenge Problems affect just about any proposed solution to the semantic paradoxes, and the most prevalent response is just to deny that the concepts in question are, in fact, expressible in the object language, either by invoking a hierarchy of metalanguages, as in Tarski (1933), or by denying that the notions are expressible at all. (Cook, 2007, pp. 33–34)

We can either view Revenge as a separate phenomenon or as the same problem as the Liar. If Revenge is a separate phenomenon then the existence of Revenge problems does not undermine a solution to the Liar paradox (but it does still seem like there is an onus on a theory to give a response to Revenge problems). If we view Revenge as the same problem as the Liar paradox, then a solution to the Liar is not complete without a response to Revenge. Either way, once we have generated a Revenge problem for a given solution to the Liar we want that theory of truth to be able to give a response to Revenge in order to consider the theory adequate.

In the next section, I will look at possible solutions to Revenge for the Liar which will be used in sections 4.4 and 4.5 to inform similar strategies that can be used for Revenge for the Probabilistic Liar.

4.2.1 Solutions to Revenge for the Liar

There have been a variety of solutions offered to the Revenge paradox. In this section, I will argue against a general style of approach: taking hierarchies of truth predicates as part of a solution to the Liar and Revenge. This will inform my approach to how we can give a solution to Revenge for the Probabilistic Liar. The idea of a hierarchy is to stratify the sentences into levels. Tarski's theory of truth involves giving a hierarchy of languages and metalanguages. It avoids the Liar Paradox by restricting the language, so there is no general truth predicate and no language can contain its own truth predicate. If you want to describe a sentence in the language $\mathcal{L}_n$ as true you have to 'step up' to the expanded language $\mathcal{L}_{n+1}$ which contains a true predicate T_{n+1} which only applies to sentences of $\mathcal{L}_m$, $m \leq n$.

Hierarchies have been offered either explicitly or implicitly as part of a number of solutions to the Liar Paradox and Revenge for the Liar. They have also been considered problematic in theories of truth since they restrict the expressive power of the language in some way. This may seem to go against our intuitions about the word 'true'. While I will argue that there are some problems with the use of hierarchies of truth predicates this is not particularly important to my position overall. I think that the use of hierarchies of truth predicates are more plausible than hierarchies of belief predicates. I will use this section to outline how hierarchies can arise in responses to Revenge. The range of responses to Revenge is too large to give an overview of all accounts that include hierarchies, but I will consider two that have different approaches. Kripke (1975) which aims to avoid explicit hierarchies but needs to invoke them to respond to Revenge, and Cook (2007) who claims we should embrace hierarchies.

This suggests that similarly, when considering responses for the Revenge problem for the Probabilistic Liar, we should avoid strategies that involve hierarchies.

Why hierarchies look problematic

Tarski's hierarchy avoids the Liar paradox but to do so he restricts the language and this is in conflict with our intuitions that the word 'true' can be applied universally rather than with implicit reference to a level of the hierarchy. The lack of a general truth predicate for the language blocks us even being able to informally state the Liar paradox. Kripke also notes that the Tarski hierarchy causes problems for our ordinary assertions about truth and falsity. He gives the example statements made by Jones and Nixon.

> Consider the ordinary statement, made by Jones:
>
> (1) Most (i.e., a majority) of Nixon's assertions about Watergate are false.
>
> [...]
>
> Suppose, however, that Nixon's assertions about Watergate are evenly balanced between the true and the false, except for one problematic case,
>
> (2) Everything Jones says about Watergate is true.
>
> Suppose in addition, that (1) is Jones's sole assertion about

Watergate. (Kripke, 1975, p.691)

In this example, Jones does not know if Nixon makes assertions that involve truth and if so at what level these assertions are made at. This suggests that sentences do not have a pre-ascribed level in the hierarchy but rather the level depends on the empirical facts (the level at which Nixon's assertions are made at). The example also illustrates that paradoxes can emerge even when there is nothing intrinsically problematic about the individual assertions. In the Tarski hierarchy (1) would have to be on a higher level than all of Nixon's utterances and (2) would have to be on a higher level than all of Jones's utterances. However, since (1) is uttered by Jones and (2) by Nixon this means (1) has to be on a higher level to (2) and (2) on a higher level to (1) which is not possible.

So, we can see that having explicit hierarchies looks problematic. This suggests that an adequate solution to the Liar paradox will need to avoid introducing a hierarchy of truth predicates.

Kripke

As I mentioned in Chapter 1 there are theories of truth that offer a solution to the Liar paradox in terms of truth-value gaps. Sentences that are paradoxical fall into these truth-value gaps. These approaches can address the Liar paradox without introducing a hierarchy of truth predicates. Kripke's fixed point theorem is a theory of truth that goes down this route. In the following, I outline Kripke's theory of truth and show how the ghost of the Tarskian hierarchy emerges and discuss whether this hierarchy is problematic.

The idea behind the fixed point approach is that truth is defined for a language (i.e. it is not taken as primitive). We start with a language L which does not contain a truth predicate and extend it to a metalanguage $\mathscr{L}$ which contains the predicate $T(x)$. The interpretation of $T(x)$ is given by (S_1, S_2) where S_1 is the extension of $T(x)$ and S_2 the anti-extension. $\mathscr{L}(S_1, S_2)$ is the interpretation of $\mathscr{L}$ as a result of interpreting $T(x)$ by (S_1, S_2). Now set S'_1 to be the set of codes of true sentences of $\mathscr{L}(S_1, S_2)$ and S'_2 to be the set of all elements of the domain which are either not codes of sentences of $\mathscr{L}(S_1, S_2)$ or are codes of false sentences of $\mathscr{L}(S_1, S_2)$. If $S_1 = S'_1$ and $S_2 = S'_2$ then we have a fixed point (i.e. this is what must be the case in order for T(x) to be interpreted as truth for L). For a given choice of (S_1, S_2) set $\phi(S_1, S_2) = (S'_1, S'_2)$, the fixed points of (S_1, S_2) are the same as the fixed points of the function ϕ i.e. when $\phi(S_1, S_2) = (S_1, S_2)$.[2]

In order to construct the fixed points Kripke asks us to consider a series of languages: $\mathscr{L}_0, \mathscr{L}_1, \mathscr{L}_2, ... \mathscr{L}_n, ...$, i.e. a ""hierarchy of languages" ...analogous to the Tarski hierarchy " (ibid, p. 703). The idea of the hierarchy being that the predicate $T(x)$ which is given for an arbitrary level of the hierarchy $\mathscr{L}_{\alpha+1}$ is interpreted as the truth predicate for the

[2] Kripke presupposes the standard device of assigned a number to each well formed formula in the language L and refers to this as the code. A standard coding device is Gödel numbering but there are also other coding devices (see (Kripke, 1975, p.702)). Bolander (2017) gives a non-formal explanation of Gödel numbering (also referred to as Gödel coding) — one can think of Gödel numbering as a way of naming or quoting a formula, so we can think of it like quotation in natural language. For a more technical explanation of how Gödel numbering works see the supplement to Raatikainen (2021) where he gives the example of how Gödel numbering can be applied to the language of arithmetic. In particular Gödel numbering gives us a way of assigning a unique number (or code) to a formula (or sentence). Kripke talks about codes of sentences "to remind the reader that syntax may be represented in L by Gödel numbering"(Kripke, 1975, p.702 fn 22).

prior level of the hierarchy $\mathscr{L}_\alpha$. The first level of the hierarchy is given by $\mathscr{L}_0 = \mathscr{L}(\emptyset,\emptyset)$, in $\mathscr{L}_0$ $T(x)$ is undefined. For any $\alpha \in \mathbb{N}, \mathscr{L}_\alpha = (S_1, S_2)$ and set $\mathscr{L}_{\alpha+1} = (S'_1, S'_2)$. It is important to note some features of the function ϕ. It is monotonic, meaning that when we extend the extension and anti-extension of T(x) this does not change the truth value of anything that had previously been defined. This means that each level of the hierarchy extends the interpretation of $T(x)$, i.e. the extension and anti-extension of T(x) increases as α increases.

'Increases' here does *not* mean 'strictly increase'. Indeed, at some point, the sentences of $\mathscr{L}$ will have been exhausted and no new sentences are declared as true or false. This gives us the minimal fixed point (any other fixed point extends this one- if a sentence is true (or false) at the minimal fixed point it will be true (or false) in any other fixed point). The intuition behind the approach is as follows: an agent starts at $\mathscr{L}_0$ where $T(x)$ is not defined, at this stage she has no understanding of truth. The agent can progress to $\mathscr{L}_1$ by being given a characterisation of truth by some scheme (i.e. the Strong Kleene valuation rules). With these rules, she can evaluate sentences that do not contain $T(x)$ as true or false. She can then extend her understanding of $T(x)$ and then evaluate more sentences as in the extension or anti-extension of $T(x)$. Eventually, she reaches a fixed point $\mathscr{L}_\sigma$, i.e. a language that contains its own truth predicate. Kripke suggests the following: a sentence is paradoxical if it does not have a truth value at any fixed point.

Kripke's fixed point approach appears to solve many of the problems that were put to Tarski's theory. The fixed point approach yields languages that contain their own truth predicates and conform to the Tarski T-biconditionals. However, with Kripke's theory, we cannot assert in the object language that the Liar sentence is not true. Kripke's solution to this is to move to a metalanguage. He says:

> … it is certainly reasonable to suppose that it is really the metalanguage predicate which expresses the "genuine" concept of truth for the closed-off language … So we still cannot avoid the need for a metalangauge … (Kripke, 1975, p.715)

It is here that we see the ghost of the Tarski hierarchy is still with us with the need to ascend to a metalanguage. Furthermore, the need to ascend to a metalanguage to assert that the Liar sentence is not true allows us to generate a Revenge problem and hierarchy from this. Let us call this metalanguage predicate 'True*'. Then we can generate a Revenge problem with True*. With our "genuine" concept of truth we can form a "genuine" Liar sentence as: 'This sentence is not True*". We then need to give an analysis of the 'genuine' Liar sentence (and it raises the question of why we were interested in the original Liar sentence at all if this is the 'genuine' form of the Liar sentence). As Maudlin puts it "We have not just 'the ghost of the Tarski hierarchy', but the full-fledged complete hierarchy, forced on us by Revenge" (Maudlin, 2007, p.193). So, if one found the Tarski hierarchy problematic then it seems Kripke's theory of truth faces the same or similar problems.

Cook

Rather than avoiding hierarchies Cook (2007) argues that with each new Revenge sentence we encounter we add a truth value: we extend the language and then when we are forced to extend the semantics as we ascend to another level and obtain a new language. This is a familiar pattern of dealing with the hierarchies but Cook argues that rather than this hierarchy approach being problematic we should instead embrace it.

> The Revenge Problem, it turns out, is not a problem at all, but affords crucial insight that allows for a truly satisfactory solution to the semantic paradoxes. (Cook, 2007, p.34)

Cook takes the fact that we seem stuck with Revenge to show that language is indefinitely extendable.

Cook gives the familiar story of each strengthened Liar paradox leading us to extend the language and thus extend the semantics. His informal justification for the position (which he later provides a formal theory for) is that truth values are "a means for keeping track of various relationships that hold between a sentence and the world" (Cook, 2007, p.37). He goes on to further explain the relationship between complexity and the number of truth values we will need:

> If truth-values are the result of different sorts of relationships that sentences can have to the world, then the number of different sorts of relationships will co-vary with two things: the complexity of the sentences, and the complexity of the world. (ibid, p.37)

So, each time we add semantic complexity by adding a new concept and thus vocabulary into our language we increase the complexity of our language (and of the world since arguably language is part of the world). It, therefore, looks natural that this will result in extending the truth-values further.

> The Revenge Problem shows us that, given any language L, L cannot contain the expressive resources necessary to completely describe the semantics appropriate for L. This is compatible with there being, for any language L, and extension of L, L', which is expressively rich enough to describe the semantics of L. What is required is an infinite sequence of semantics, each one expressively richer than the last, and an infinite sequence of semantics, one for each language. (ibid, p.39)

According to Cook the Revenge problem is in fact not a problem at all but rather highlights a feature of the language.

The real worry for Cook's approach comes by questioning why we cannot just quantify over all the added truth values to generate a super strengthened Revenge problem. As with any explicit hierarchy approach this account needs to give a principled reason for not allowing universal quantification.

Cook gives some arguments for why this should not be problematic for his account. He points out that the idea of not being able to quantify over indefinitely extensible concepts is familiar from set-theoretic paradoxes and the inability to quantify over all sets.

> ...there are some properties such that, given any class of terms all having such a property, we can always define a new term also having the property in question. Hence we can never collect all the terms having the said property into a whole; because, whenever we hope we have them all, the collection which we have immediately proceeds to generate a new term also having the said property. (Russell, 1906, p.36)

While there are some compelling features of this account, I will argue in section 4.4 that this approach looks problematic when we consider it in a belief context. His approach also still stratifies truth and has restrictions on expressive power that come with accepting hierarchies.

4.3 Revenge for the Probabilistic Liar

The Probabilistic Liar has clear parallels to the Liar paradox. In the previous section, I outlined how a Revenge problem can be given for the Liar paradox if one gives a solution to the Liar that introduces a new semantic category. Once we have introduced a new category to resolve the original problem we leave ourselves open to generating a Revenge problem in terms of that new semantic category. In this section, I will show how we can similarly generate a Revenge problem for the Probabilistic Liar in terms of the solution I have argued for in the previous chapters.

Recall, the version of the Probabilistic Liar that I am using can be formed by imagining an agent Alex who stipulates the meaning of a sentence (α) to be the following:

(α) Alex's credence in the proposition expressed by (α) is less than 0.5

In more concise notation:

(α) $Cr\rho(\alpha) < 0.5$

In the previous chapters, I have argued that an agent should suspend judgment in the Probabilistic Liar. This solution has a clear parallel to a solution given to the Liar paradox - the solution of adding a new semantic category. We have seen that the Probabilistic Liar is problematic if we assume that the only determinate attitudes the agent can take towards it are precise credences. This goes away however once we consider the possibility of having imprecise credences and in particular credence gaps (or suspending judgement). This suggests that we can similarly generate a Revenge problem for the Probabilistic Liar.

If we can express a new proposition that refers to an agent's attitude being suspended judgment, then we can form a strengthened Probabilistic Liar. That is if we formalise the following:

(β) : Alex either suspends judgement in proposition expressed by (β) or has credence less than 0.5 in (β).

then we have a Revenge problem. We can see that, as with the Probabilistic Liar, (β) is problematic. The intuitive problem can easily be given: if Alex suspends judgment in (β) then (β) is true (and so Alex should have some precise high credence in (β)). However, if Alex has any precise credence in (β) then we have the same problem as the Probabilistic Liar and thus one should suspend judgment in (β). If we can formalise the Revenge problem then it seems that we are back with the original problem of the Probabilistic Liar. Once we include the types of attitudes an agent can have we can generate new problematic examples of probabilistic self-reference.

As we saw in the possible ways one can respond to Revenge problems for the Liar, one way we can respond to Revenge is to deny that this is expressible at all. This does not seem like a fruitful strategy, since, if we consider the formal representation of suspended judgement we can see how a Revenge problem for the Probabilistic Liar looks easily expressible. A formal representation of Revenge also makes clear the possible ways we can respond to the problem.

I have argued that we should understand suspended judgment as having imprecise credences and in particular having a credence gap. Formally a Revenge problem for the Probabilistic Liar will emerge if we can talk about having imprecise credences or refer to the determinate attitudes an agent may have. Consider the following:

$$(\beta) \quad C\rho(\beta) < 0.5$$

Where $C\rho(\beta)$ is the credal set Alex adopts towards (β).

The intuitive problem with (β) looks much the same as the intuitive problem for (α). If $C\rho(\beta) < 0.5$ then it follows that (β) is true and so, in light of this it looks like Alex ought to in fact be confident in (β), i.e $C\rho(\beta) > 0.5$. If however, all the credence functions in her credal set assign (β) credence greater than 0.5 then (β) is false so it looks like she ought to have low credence in (β), i.e $C\rho(\beta) < 0.5$. Note there is another possibility to consider, when $\neg C\rho(\beta) < 0.5$, in this case all that is required is that at least one credence function in $C\rho(\beta) \geq 0.5$ rather than all the credence functions in the credal set. In this case it again follows that (β) is false and so, in light of this she ought to have $C\rho(\beta) < 0.5$.

So, once again, as with the Probabilistic Liar, it looks like there is no stable attitude an agent can adopt towards (β). Moreover, we have the requirement that all the credence functions in the credal set obey Probabilism. Given this, if we can also give similar versions of the positive and negative introspection principles Caie gives for precise credences, then it looks like the same argument as Caie gave for the Probabilistic Liar emerges. I discuss the introspection principles and argument in Appendix A.

When faced with this problem for the Probabilistic Liar the strategy was to move to suspended judgment - i.e. have a credal set rather than a precise credence in the proposition. We can see however that suspending judgment does not help us here. If Alex has a determinate attitude to (β) represented by a credal set then it is determinate whether $C\rho(\beta) < 0.5$ or not. As I noted above, if Alex has any determinate attitude then, given that attitude she will be certain whether (β) is true or false. In light of this, she ought to change her attitude towards (β). Even if Alex suspends judgment she has a determinate attitude and if she can introspect on her determinate attitudes then a Revenge problem can be generated.

As such, it looks like once we consider an agent's determinate attitudes being represented by an imprecise credal set we can express a

Revenge problem. Given the expressibility of a Revenge problem for the Probabilistic Liar in the next sections I will outline how we might respond.

4.4 Possible Solutions to Revenge for the Probabilistic Liar

As I discussed in section 4.2.1 there are a number of different solutions to the Liar and Revenge. I argued that accounts that use hierarchies do not provide good solutions to Revenge. Similarly, I think that analogous solutions to Revenge for the Probabilistic Liar will not be fruitful strategies. In this section I will consider analogous solutions for Revenge for the Probabilistic Liar and show that they will also not act as adequate solutions.

A Tarski style solution would involve the Probabilistic Liar not being expressible. By removing the possibility of probabilistic self-reference from the language we can block the original Probabilistic Liar and thus the possibility of Revenge as well. In a similar style to Tarski, we could introduce a hierarchy of languages. In L_0 we cannot talk about probabilities at all, in L_1 we can talk about probabilities of sentences in L_0, etc. Just as in the alethic case this would mean there is a lack of expressive power. It seems that in natural language self-referential probability can occur and we, therefore, ought to be able to express this in a formal language. Campbell-Moore (2015a) gives the following example:

> Suppose that Smith is a Prime Ministerial candidate and

> the candidates are campaigning hard today. Smith might say:
>
> (i) "I don't have high credence in anything that the man who will be Prime Minister says today."
>
> Imagine further, that unknown to Smith, he himself will become Prime Minister. (i) therefore expresses a self-referential probability assertion. (Campbell-Moore, 2015a, p.683)

Removing the possibility of expressing self-referential probability from a formal language, therefore, does not look like a plausible approach.[3]

If we accept self-referential probability into the formal language we can then look at how Revenge impacts that solution. With both Kripke (1975) and Cook (2007) their solutions involve invoking a hierarchy of truth predicates. In Kripke's case, he has to ascend to the meta-language and the 'genuine' notion of truth. Maudlin (2007) argues that this ends up meaning Kripke needs the Tarskian hierarchy to accommodate Revenge. Cook (2007) embraces expanding the language to introduce a new pathological category.

In both of these cases, the analogous move that it seems we would need to make in the belief case is to either introduce a new 'genuine' notion of credence or introduce a hierarchy of suspended judgment attitudes. Looking at each of these options in turn it seems that neither of these looks like promising approaches to Revenge for the Probabilistic Liar.

[3] See (Campbell-Moore, 2015a, §2) for future examples.

Let us consider introducing a new sense in which one could believe or disbelieve the Probabilistic Liar. Let the 'genuine' credence attitude credence* be the genuine notion of having an attitude towards a proposition. This suggests we can reformulate the probabilistic Liar (α) in terms of Alex having some credence* towards (α).

$$(\alpha) \quad C^*(\alpha) < 0.5$$

and this would be the 'genuine' Probabilistic Liar. This raises the question of why we were interested in C when C^* was the genuine attitude and we can see that just as for Kripke, a hierarchy will re-emerge if we repeat this reasoning.

Looking at a Cook style solution to Revenge for the Probabilistic Liar we could instead embrace hierarchies of Suspended Judgment, Suspended Judgment$_2$, Suspended Judgment$_3$, Suspended Judgment$_4$, …etc. While we might be able to make sense of a hierarchy of semantic categories the belief case is hard to make sense of. It is not clear what these different attitudes would be or how to understand them.

In my argument for suspended judgment as the attitude an agent ought to adopt towards the Probabilistic Liar I give an account of how this attitude should be understood. Suspended judgment is a neutral attitude an agent can have towards a proposition. It is also a natural third category to consider when we look at it in the context of a traditional belief framework. Given this, it is also natural to consider how it can be expressed in a credal framework. It is unclear how we could make sense of further distinct attitudes that an agent could have towards a proposition while also not having any precise credences nor suspending judgment. All of this points to finding an approach that avoids hierarchies of different belief attitudes. In the next section I will suggest a solution that avoids hierarchies.

4.5 A solution to Revenge for the Probabilistic Liar

As we saw with various responses to Revenge for the Liar that I gave in section 4.2, the types of response that are offered depend on certain underlying assumptions about expressive power and assertability. In this section, I will examine the assumptions that are needed to generate the Revenge problem for the Probabilistic Liar and use these to suggest a solution.

In the Revenge problem for the Probabilistic Liar we can see that there seem to be two key assumptions that generate the problem:

a) That Alex has a determinate attitude towards (β), so there is some determinate set of credence functions that represents her attitude towards (β).

b) That Alex can introspect on her determinate attitudes and therefore be confident whether $C\rho(\beta) < 0.5$ or not.

That is to say, if the agent is aware that she has a credence gap towards a proposition then we can give a Revenge problem where suspending judgment will not work as a solution. Suspended judgment is a determinate attitude the agent has towards a proposition, as such if Alex has an intersectionist credal set $C = \{Cr_1, Cr_2, Cr_3, \dots\}$ then she has a determinate attitude towards (β) and thus it is determinate whether $C\rho(\beta) < 0.5$ or not.

Given these assumptions it, therefore, suggests two strategies we could employ to respond to the Revenge problem:

1) Find a way for it to be indeterminate what attitude Alex has towards (β).

2) Limit Alex's introspection so she is unsure of whether she has a credence gap towards a proposition.

Strategy 2) of simply restricting Alex's introspection so she is not confident of her determinate attitude towards a proposition seems unmotivated. As in the original Probabilistic Liar argument from Caie that I gave in Chapter 1, we can see that, although it may be reasonable for an agent to sometimes have limited introspection, it seems problematic that an agent is rationally required to have poor introspection on her own determinate attitudes. This suggests looking at strategy 1) and how we can represent that it is indeterminate what attitude Alex has.

In my proposed solution to the Probabilistic Liar, the solution questioned the assumption that an agent must have complete precise credences. Formally the solution involved going imprecise and representing an agent's determinate attitude in terms of a set of credence functions. This suggests a way of expressing a distinct attitude an agent can take towards (β). We can question the assumption that we can represent an agent's attitudes in terms of a single credal set. Instead, we can consider there being a set of credal sets and this is the attitude Alex takes towards (β). This formal move does not yet provide a solution to the Revenge problem, we need an explanation of the interpretation of the set of credal sets and how this interpretation blocks the problematic reasoning. As I discussed in the previous chapter, one prominent interpretation of imprecise credences is the supervaluationist interpretation. With this interpretation we can represent that it is indeterminate what an agent's attitude is. That is, by taking a set of credal sets we can take strategy 1). It is indeterminate what attitude Alex has towards (β).

In the following section I will argue that the supervaluationist interpretation makes sense as an interpretation of a set of credal sets. Using this I will then revisit the proposed strategy for a response to Revenge and show that taking a set of credal sets and interpreting this with a supervaluationist interpretation acts as a solution.

4.5.1 Supervaluationism

In the previous chapter I outlined how the supervaluationist interpretation can be applied to a single credal set. In this case, we interpret the credal set as representing that each of the precise credence functions in the credal set is a permissible precisification of the agent's attitude. However, overall the credal set represents that it is indeterminate what the agent's attitude is. I have argued however that we ought to represent an agent's comparative beliefs with a set of credence functions rather than with a precise credence function. We can therefore also view the supervaluationist interpretation as applying to an agent's determinate comparative beliefs when these determinate comparative beliefs are represented by a credal set. It can be indeterminate or vague which comparative attitudes an agent holds.

If we take the determinate attitudes of an agent to be represented by credal sets C (which we interpret using an intersectionist interpretation) then by taking a set of these we can represent that it is indeterminate which intersectionist credal set C represents the agent's attitude. Let C_* denote a set of credal sets and each C_i in the set represents an intersectionist credal set:

$$C_* = \{C_1, C_2, C_3, \ldots\}$$

The credal set C_* represents that it is indeterminate which of the C_i in the set represent the agent's belief.

This shows that the supervaluationist interpretation of a set of credal sets is compatible with also having an intersectionist interpretation of imprecise credences. So, there is no worry that adopting a supervaluationist interpretation for a set of credal sets undermines my argument for an intersectionist interpretation. The intersectionist interpretation naturally follows from the comparativist framework which I have argued for, and expresses an agent's determinate attitudes. It is also possible that an agent with comparative beliefs has indeterminacy or vagueness in their attitudes as well. This indeterminacy can be represented with a supervaluationist interpretation applied to sets of credal sets.

Using the supervaluationist interpretation we can therefore represent that it is indeterminate what Alex's attitude to (β) is and give a solution to Revenge for the Probabilistic Liar. Let Alex's set of credal sets in (β) be as follows.

$$C_*\rho(\beta) = \{C_1, C_2, C_3, \ldots, C_n\}$$

Where each of the C_i are credal sets and $C_i\,\rho(\beta) < 0.5, \neg C_j\,\rho(\beta) < 0.5$ for some $C_i, C_j \in C_*$.

Moreover, it looks like, when it is indeterminate what an agent's attitudes are, we can give a principled weakening of introspection principles. It seems plausible that an agent can introspect on her attitudes and be confident of her determinate attitudes. However, when it is indeterminate what an agent's attitude towards (β) is it is unclear why we should expect an agent to be confident about what credal sets represent her attitude.

However, as we saw with Revenge for the Liar, Revenge appears to be a pervasive problem for an account. Once we have offered a solution to a Revenge problem we can once again ask the question of whether a new strengthened Revenge problem can be generated in terms of this solution. The solution I have offered for the Revenge problem for the Probabilistic Liar is based on being able to express that it can be indeterminate what an agent's attitudes is as well as expressing their determinate attitudes. This indeterminacy can be formally represented by a set of credal sets. If we can refer to an agent's set of credal sets then it seems we once again generate a new problematic sentence.

We can generate a new problematic sentence by talking about the set of credal sets C_*. In particular, if we can refer to the set of credal sets as representing that it is indeterminate what an agent's attitudes are. Informally:

(γ) It is indeterminate whether there is a credal set $C_i \in C_*\rho(\gamma)$ such that $C_i\,\rho(\gamma) < 0.5$.

The solution of taking a set of credal sets will not work for (γ) since (γ) is self-referential with respect to that level of set complexity. Alex's set of credal sets represents that it is indeterminate whether there is a credal set $C_i \in C_{*}\rho(\gamma)$ s.t $C_i\, \rho(\gamma) < 0.5$. This makes (γ) true, so Alex ought to in fact be confident in (γ). If she is confident in (γ) if follows that she has a determinate attitude, so it is not indeterminate whether there is a credal set $C_i \in C_{*}\rho(\gamma)$ is s.t $C_i\, \rho(\gamma) < 0.5$. In which case (γ) is false, so she ought to in fact be unconfident in (γ). Again, when she is unconfident in (γ) it follows that she has a determinate attitude towards (γ). So there is no stable attitude she can take towards (γ).

Following the same reasoning as before we can see that to resolve this new problem, we cannot appeal to it being indeterminate what credal set represents her attitude. Since (γ) refers to the indeterminacy of $C_{*}\rho(\gamma)$ we cannot appeal to $C_{*}\rho(\gamma)$ representing the indeterminacy of Alex's attitude. Just as before we can consider the case where it is indeterminate which set of credal sets she takes as permissible precisifications of her attitude towards (γ). This can be represented by the agent having a set of sets of credal sets. Again, since we want to represent indeterminacy in the agent's attitude we give a supervaluationist interpretation of the set of sets of credal sets. The set of sets of credal sets represents that it is indeterminate which of the sets of credal sets represents the agent's attitudes. Indeed, if an agent's attitude are represented by some iteration of credal sets then it looks like we can simply keep taking sets of sets of credal sets or sets of sets of sets of credal sets to describe an agent has indeterminacy in her beliefs about what credal sets represent her attitude.

Once again we have given a solution to a Revenge problem by adding in a layer of set complexity. This suggests that we will continue to be able to generate new Revenge problems and generate solutions that will give us an iteration of sets of credal sets. One might worry that, with this approach that keeps adding in levels of set complexity, we also need to have a philosophical story or justification for why we can do this. Does these make sense as attitudes agents could take? I have argued in Chapters 2 and 3 that we can make sense of an agent's determinate attitudes in terms of a credal set with an intersectionist interpretation. This follows if we accept that comparative beliefs are what is psychologically real and the set of credences is taken as a way of representing that an agent can have incomplete comparative belief orderings. In the previous section I also argued that as well as having determinate attitudes it might be indeterminate what an agent's attitudes are. We can represent this indeterminacy with a supervaluationist interpretation of a credal set. If an agent's determinate attitudes are represented by (intersectionist) credal sets then we can make sense of a set of credal sets. We interpret it as it being indeterminate which credal set represents the agent's attitude.

We might also think that we can make sense of it being indeterminate which set of credal sets represents an agent's beliefs. Again, as I noted in the previous section, this can be represented with a supervaluationist interpretation. It seems however that after a certain level of set complexity it is very unrealistic that an agent could have this level of complexity their attitude. For real agents anything above a set of sets of credal sets looks like it may already be too complex to conceive of. If the attitude that I am recommending an agent has in response to each successive Revenge problem is unrealistic does this undermine my response to Revenge?

The unrealistic complexity of the attitude is not a problem. Once we get past a certain level of complexity the strengthened Revenge problems also become very complex to formulate. However, since they are possible to formulate, an account ought to be able to explain what attitude an ideal agent ought to theoretically have towards them. It merely needs to be possible that an ideal agent could have these attitudes. Since the supervaluationist interpretation can make sense of sets of credal sets, sets of sets of credal sets and so on, this is all that is needed for the possibility that it could be indeterminate what the agent's attitude is at these levels of set complexity. Moreover, it is not obvious whether the complexity of the set structure requires extra cognitive complexity. When we consider what attitude an agent ought to have towards 'Patchy is red', where Patchy is a borderline colour patch, it looks like their attitude ought to be indeterminate. It follows that any indeterminacy that is caused by vagueness will already generate indeterminacy in belief. Rather than requiring cognitive complexity, it seems like all we require is that it is vague what attitudes the agent has.[4]

This shows that we can make sense of further Revenge problems and how we can respond to them. The method of generating further Revenge problems suggests that Revenge problems form a hierarchy and as I discussed in section 4.2.1 hierarchies seem problematic. The

[4] I will discuss the types of attitude an agent ought to take towards indeterminacy in more depth in the next chapter.

attitude in each case is the same however. At each level of set complexity it is indeterminate what attitude the agent has in response to the higher-order indeterminacy in the formulation of the Revenge problem. In the next section I will show how Revenge can be formed in terms of an iteration of determinacy operators. The approach presented above does not seem to necessitate a problematic hierarchy like those I mentioned in section 4.2.1.[5]

One worry that does emerge is whether a super-Revenge problem might be generated depending on the logic of determinacy one accepts.

4.6 Revenge in terms of indeterminacy

In the previous section I formulated Revenge in terms of the credal set structure. However, we can see that further Revenge problems for the Probabilistic Liar can be given in terms of the lack of determinacy of an agent's attitude. Note we can define $\triangledown p$ to mean it is indeterminate that p and we can express indeterminacy in terms of determinacy as follows:

$$\triangledown p = \neg \vartriangle p \wedge \neg \vartriangle \neg p$$

If we define a 'maybe' operator as $\neg\vartriangle\neg$ then we can formulate Revenge and the attitudes I recommend as solutions in terms of a determinacy operator.

[5]Talking about that determinacy operator does however require a theory of vagueness. As I will discuss in section 4.6.2 once we refer to the determinacy of an agent's attitudes in the formulation of the Revenge problem we also have to address a question about the logic of determinacy.

$$(\beta)\ \ C\rho(\beta) < 0.5$$

As we saw, if it is determinate whether $C\rho(\beta) < 0.5$ then this proposition looks problematic. I suggested that the attitude an agent should adopt towards this was to have a set of sets of credences and thus it is indeterminate whether $C\rho(\beta) < 0.5$. From this it follows:

$$\neg \vartriangle \neg\ C\rho(\beta) < 0.5$$

This then suggests that another strengthened Revenge could be generated. In terms of the maybe operator we can write this as follows:

$$(\gamma)\ \ \neg \vartriangle \neg\ C\rho(\gamma) < 0.5$$

That is to say "maybe $C\rho(\gamma) < 0.5$". As we can see, we cannot use the solution suggested above of taking a set of credal sets where in some credal sets it is the case that $C\rho(\gamma) < 0.5$ and in others not. It would then be the case that $\neg \vartriangle \neg\ C\rho(\gamma) < 0.5$. Instead, the agent should have a set of set of sets of credal sets where it is borderline if 'maybe' the case one of these credal sets expresses $C\rho(\gamma) < 0.5$. From this it follows:

$$(\neg \vartriangle \neg)(\neg \vartriangle \neg)\ C\rho(\gamma) < 0.5$$

Given this solution we can easily give another strengthened Revenge problem where this approach will not suffice. We can give this by applying the 'maybe' operator again to get that it's 'maybe' 'maybe' that

$C\rho(\pi) < 0.5$.

$$(\pi) \;\; \neg \vartriangle \neg\neg \vartriangle \neg \; C\rho(\pi) < 0.5$$

or, simplifying this:

$$(\pi) \;\; \neg \vartriangle \vartriangle \neg \; C\rho(\pi) < 0.5$$

Each time a strengthened Revenge can be given by again applying another 'maybe' operator. We can give a solution by giving another layer of set complexity and it being borderline if 'maybe' 'maybe' $C\rho(\pi) <$ 0.5.

In general we can give the n^{th} Revenge problem as:

$$(\sigma_n) \;\; \neg \overbrace{\vartriangle \cdots \vartriangle}^{n-1} \neg \;\; C\rho(\sigma_n) < 0.5$$

for any $n \in \mathbb{N}$.

Earlier I argued against a hierarchy solution to Revenge for the Probabilistic Liar. One might worry that with my current solution we generate a hierarchy of determinacy operators. My solution generates an iteration of determinacy operators rather than a hierarchy. The attitudes that I recommend an agent adopt to each new strengthened Revenge problem do not form a hierarchy of different attitudes, rather it is the same attitude an agent has in each case- it is indeterminate what attitude they have. This indeterminacy can be applied at different levels of set complexity but it is ultimately the same attitude. Expressing the Revenge problem in terms of a 'maybe' operator we can see then that the type of solution I have given embraces an iteration of the determinacy operators.

From a first look at the iterated maybe operators as solutions to the various strengthened Revenge problems we can see that, given certain logics, we can simplify the 'solutions'. Given the modal operators, we can understand determinacy formally by analogy to the possible worlds structure (rather than a set of possible worlds we take a set of points that correspond to possible sharpenings). Just as we can define what it means for $\Box A$ to be true we can give a similar definition for $\triangle$. We say 'determinately A' is true at some point x if and only if A is true at every point x admits.[6] An overview of axioms of modal logic that will be relevant are as follows. Let us start with a weak modal logic on which the following principles hold:[7]

(Closure): $\triangle(A \rightarrow B) \rightarrow (\triangle A \rightarrow \triangle B)$

(Necessitation): If A is a theorem of the logic then so is $\triangle A$

Other axioms can be added depending on the frame relation. If the frame relation is reflexive, we also have the axiom **T**:

T: $\triangle A \rightarrow A$

If the frame relation transitive we get the axiom **4**:

4: $\triangle A \rightarrow \triangle \triangle A$

and if the frame relation is both reflexive and transitive we get the logic **S4**, i.e **S4** is the system that results from adding **4** to **T**.

If the frame relation is symmetric we can add the axiom **B**:

[6] The admitting relation between points is analogous to the accessibility relation between possible worlds.

[7] Garson (2021).

B: $A \to \triangle\neg\triangle\neg A$

When **B** is added to **S4** we get the logic **S5**. We can also get the logic **S5** from adding **T** and the following axiom:

5: $\neg\triangle\neg A \to \triangle\neg\triangle\neg A$

In the rest of the chapter I will assume that the frame relation is reflexive and therefore the **T** axiom holds. Given this, when I talk about the **4** axiom it will follow that we have the **S4** logic and, when I talk about the **5** axiom it will follow that we have the **S5** logic.

Now, if we take the solution to (γ):

$$\neg\triangle\neg\neg\triangle\neg C\rho(\gamma) < 0.5$$

By double negation elimination it follows that:

$$\neg\triangle\triangle\neg C\rho(\gamma) < 0.5$$

By either **S4** or **S5** it follows that we can contract strings of $\triangle\cdots\triangle$ to $\triangle$ (see (Garson, 2021, §2)). So it follows that:

$$\neg\triangle\neg C\rho(\gamma) < 0.5$$

Similarly for the solution to (π) (the strengthened Revenge problem). So, in order to have a solution the logic of the determinacy operator must be weaker than **S4** or **S5** (since we could simplify the attitude to get paradox).

Moreover, in order for first-order Revenge to work it need to be possible for it to be indeterminate or vague whether something is determinately p. The claim that it is never vague that it is determinately p can be expressed as follows:

$$\triangle \triangle p \vee \triangle \neg \triangle p \tag{4.1}$$

If (4.1) is true then it bans indeterminacy of whether or not it is determinate that one's attitude to (γ) is less than 0.5. This rules out my response to Revenge. So, it follows that in order for my solution to work we need to rule out (4.1).

Given the factivity of $\triangle$ and excluded middle (4.1) follows from:

4: $\triangle A \rightarrow \triangle \triangle A$

and

5: $\neg \triangle A \rightarrow \triangle \neg \triangle A$

i.e it follows from the **4** and **5** axioms. So for my solution to work we need to deny at least one of these claims.

Given that I'm assuming **T** holds if I reject **4** it follows that I also reject the **S4** logic. [8] Rejecting **4** seems unproblematic in the case just presented. We can deny that the accessibility relation for determinacy is transitive. This seems natural on a supervaluationist framework. When we consider the set of points, we are interested in the admitting relation between these points. In the supervaluationist in-

[8] Also note **S5** can be formulated by adding **B** to **4** and **T**. **S5** can also be formulated by adding **5** to **T**. Given that I am holding **T** fixed we cannot reject the **S4** logic and also have the **S5** logic (since the **S5** logic applies to frames that are transitive, reflexive and symmetric and the **S4** logic applies to frames that are transitive and reflexive).

terpretation, points *s* and *t* may not differ by very much and so *s* admits *t* and similarly a points *t* and *u* may not differ very much and so *t* admits *u*. However, *s* and *u* may differ enough that *s* does not admit *u*.[9]

Rejecting **S4** (and **S5**) looks reasonable in this case, but as I will show in section 4.6.2 it looks like a version of (4.1) reappears when we consider higher-order indeterminacy. This is a problem that has been presented more generally regarding whether the logic of the determinacy operator can accommodate higher-order vagueness. If higher-order vagueness can cut out, then we can express a super strengthened Revenge problem to which the above pattern of response would not apply. If it can, then this would prove problematic for my account. In the next section I will discuss higher-order indeterminacy and use this to inform a discussion on the logic of determinacy in section 4.6.2.

4.6.1 Higher-order Indeterminacy

A well known problem for supervaluationism (and many accounts of vagueness) is the problem of higher-order vagueness. Even if we can account for the initial indeterminacy there is a new problem of

[9] There are a number of ways we can understand the model here which depends on how one views the source of the indeterminacy in an agent's attitude. As Elliott (ms) notes "maybe the source of that indeterminacy is really out there in the world (or in [the agent's] head) [...] or maybe it's only a property of the language we use to talk about beliefs and their strengths, a consequence of some semantic indecision on our part." (p.6). If the indeterminacy is a property of the language then we can view the attitude the agent adopts as the permissible attitude under each precisification of the language. We can then directly adopt the same interpretation of the admitting relation and points that supervaluationism does. Namely, that we have a space of all admissible interpretations/ precisifications (the points) and the accessibility relation is "imposed upon by a given space of precisifications" (Varzi, 2007, fn9 p.639). (See also (Williamson, 1994, §5.6)).

higher-order vagueness. In this section I will highlight why an account of vagueness ought to accommodate higher-order vagueness (or indeterminacy) and how supervaluationism faces a potential problem with accommodating higher-order vagueness. This will be used to inform the discussion in section 4.6.2 where I discuss the logic of determinacy in light of the nature of the admissibility relation in a supervaluationist framework.

If we consider a vague predicate such as 'red' we typically think that there are cases that are clearly red, cases that are clearly not red and some borderline cases. That is to say, there isn't a sharp cut-off between the clearly red and the clearly not-red cases.

What goes wrong if we deny there can be borderline borderline cases? One approach is to argue that there is *vagueness* but not *higher-order vagueness*. That is, we explain why there isn't a sharp cut-off between red and not-red by the existence of borderline cases but say there is a sharp cut off between the determinately red and the borderline red cases. This seems problematic for a number of reasons. For one, just as we thought it was unrealistic that there be a sharp cut off at the first order it seems unrealistic that there be a sharp cut off at the higher order. It seems that the explanation for why ought to be the same between these two cases. At the first order, we can explain by appeal to vagueness. If we deny higher-order vagueness we either have to accept sharp cut-offs one level higher or we would have to introduce a new concept to explain the lack of sharp cut off. Bacon (2020) suggests one could introduce the notion of *schmagueness* as the thing that prevents us from knowing the sharp cut off between things that are determinately *p* and borderline *p*. We can then classify things as neither vague nor schmague. If we cannot have higher-order schmagueness we can introduce another new notion etc. By doing this we could postulate an infinite hierarchy of different notions (Bacon, 2020, p.169). [10] Denying higher-order vagueness looks like it would just give rise to a similar problem. Given this, it looks like an account of vagueness should accommodate higher-order vagueness. So, a condition on a theory should be that it at least admits all finite orders of vagueness.

Supervaluationism (in both its alethic and probabilistic forms) faces the problem of sharp boundaries. In the alethic case, borderline statements lack truth value. There are admissible interpretations, and each admissible interpretation is precise but there can be a plurality of admissible interpretations for a given term. If each admissible valuation makes a sentence true, then we can say without qualification that it is true. Similarly for false. (Which we refer to as supertrue and superfalse). A supervaluation divides the sentences of the language into three categories: true, false and neither true nor false. The problem of sharp boundaries arises here if there is a precise set of admissible interpretations of a vague term like 'red'. If the set of admissible interpretations is precise then we have a sharp higher-order boundary.

At first glance it might look problematic for the supervaluationist, if it's the case that a sentence is true on all admissible interpretations

[10] Williamson also suggests this approach saying if one wanted to insist on higher-order vagueness they could define a new operator 'definitely!' which may, in turn, require we introduce another operator 'definitely!!' and this process could continue with no natural end. (Williamson, 1994, p.160)

then it is super-true and if it is not true on all admissible interpretations it is either borderline or false. So, it seems like, given a set of admissible interpretations we have sharp boundaries and supervaluationism cannot accommodate higher-order vagueness. This can be resisted however by denying that there is some precise range of admissible interpretations.[11] If the range of admissible interpretations is not precise then the general notion of admissible will itself be vague. That is, by having the admissible interpretations be vague there is no sharp cut-off.

The problems of higher-order vagueness that I have discussed above also apply to alethic supervaluationism. As Rinard (2017) notes, problems for existing accounts of vagueness survive generalisation to the doxastic case. A parallel problem for the supervaluationist interpretation of imprecise credences is whether there are sharp cut-offs in the sets of imprecise credences. When we consider a credal set there seems to be precise endpoints to the set, which gives us sharp cut-offs.

If it is indeterminate what the agent's attitude towards 'patchy is red' is (where patchy is a borderline case of red) then, if there are sharp upper and lower probabilities this suggests that there is a sharp cut-off between the credences that are permissible precifications of an agent's attitude towards 'patchy is red' and those credences that are not permissible precifications of an agent's attitude. Just as in the case with alethic supervaluationism this gives a sharp boundary where

[11] See (Keefe, 2000, Ch.8).

it looks like there shouldn't be one. It looks like, as with the alethic case, we must deny that there are sharp cut-offs at the endpoints of the credal set.

There have been solutions offered for this problem such as the suggestion given by Rinard (2017) who suggests a weaker interpretation of supervaluationism. She calls this the minimal interpretation. If an agent has the credal set $[a, b]$ then we interpret this as meaning that:

1) Every number outside $[a, b]$ is determinately not the agent's credences.

2) The interval may contain some numbers that are determinately not the agent's credence, or even a number that determinately is the agent's credence.

On this view, there can be multiple intervals that describe your attitude. Thus avoiding any issues with sharp end-points to an interval since the interval given is not taken to be unique.[12] This points to one way the problem can be resolved.

In the next section I discuss a problem for higher-order vagueness and the possibility that, given certain logical principles, we have to say that vagueness turns out to just be a surface phenomenon.

4.6.2 The logic of determinacy

An issue arises for my proposed solution to Revenge when we consider the structure of higher-order vagueness and what this implies

[12] The downside of this view is that there is no requirement to give the narrowest interval.

for the logic of the determinacy operator. In this section I will present Williamson's problem for higher-order vagueness and how we can respond by giving good reasons to weaken the logic for the determinacy operator.

Williamson (1994) presents a problem by looking at the structure of higher-order vagueness. The intuition behind higher-order vagueness is that just as we have borderline cases of red we also have borderline borderline cases of red, borderline borderline borderline cases of red ...etc. What does a formal definition of nth order vagueness look like? We can give such a definition in terms of a determinacy operator Δ where we treat the determinacy operator as a modal operator (reading the modal necessity operator as 'determinately', and we read the modal possibility operator as 'maybe' as defined in section 4.6). Williamson gives a problem for higher-order vagueness on the supervaluationist account that emerges when we consider what frame conditions must hold for a supervaluationist interpretation of vagueness (and thus what logic holds for determinacy).

Given the modal operators, we can understand higher-order vagueness formally by analogy to the possible worlds structure. From this, we can give a formal definition of what higher-order vagueness is:

> Formally, higher-order vagueness corresponds to contingency in which worlds are possible ... For the supervaluationist, definiteness[13] under a sharpening s is truth under all sharpenings admitted (deemed admissible) by s; a

[13]The terms 'definitely' and 'determinately' are both used in the vagueness literature.

> sharpening t may be admitted by a sharpening s_1 and not by a sharpening s_2. (Williamson, 1994, p.128)

The question then, is how we should understand the accessibility relation in the vagueness framework. There are certain frame conditions that look like they must apply to the supervaluationist interpretation (which look problematic when we consider higher-order vagueness). To say A is valid we write "$\models A$". The following principles I gave in section 4.5.1 hold on Williamson's framework:

(Closure) K: $\models \vartriangle(A \rightarrow B) \rightarrow (\vartriangle A \rightarrow \vartriangle B)$

(Necessitation) RN: If $\models A$ then $\models \vartriangle A$

When we consider the admissibility relation R for supervaluationism every sharpening admits itself, hence it follows a supervaluationist interpretation has reflexivity. This gives us the axiom of

T : $\models \vartriangle A \rightarrow A$

Williamson also defends R being symmetric for the supervaluationist.

> ... a sharpening s has R to a sharpening t just in case s admits t; if sharpenings deem admissible just those sharpenings that differ from them by at most some fixed amount, then R will ...be symmetric. (Williamson, 1994, p.130)

This gives us the axiom of

(Brouwer's Principle) **B**: $\models A \rightarrow \vartriangle\neg \vartriangle \neg A$

Once we have a determinacy operator we can iterate it such as in the sequence: $\vartriangle, \vartriangle\vartriangle, \vartriangle\vartriangle\vartriangle, \dots, \vartriangle^i, \dots$. We can also define a new operator 'determinately*' A to mean the conjunction of A and determinately

A and determinately determinately A and ...etc. Or as Bacon (2020) defines it, $\vartriangle p$ means that p and it's not vague that p and $\vartriangle^* p$ means that p and it's neither vague nor higher order vague that p. We can write this as an infinite conjunction: $\vartriangle^* p := \bigwedge_{n \in \omega} \vartriangle^n p$ for some finite order n.

With determinately* defined we can ask the same question about the strength of the modal logic and if it can be vague whether something is determinately* p. As I have shown in section 4.6 we can give Revenge problems in terms of iterated determinacy operators and after each iteration, the response is that it is indeterminate what attitude the agent adopts towards that Revenge problem. It looks like by using 'determinately*' in place of 'determinately' we can generate a super-Revenge problem (unless there is a way to deny the conclusion that it is never vague whether something is determinately*).

A super-Revenge problem in terms of $\vartriangle^*$ is as follows:

$$(\pi^*) \quad \neg \vartriangle^* \neg C\rho(\pi^*) < 0.5$$

Importantly, in order to give a solution in the pattern I gave above we need to be able to express that it can be vague whether something is determinately* p. Since in order to respond to super-Revenge the right mental state for an agent to be in is one where it's indeterminate whether it's super-determinate that $C\rho(\pi^*) < 0.5$.

The claim that it is never vague whether something is determinately*

p can be given by the disjunction:

$$\triangle \triangle^{*} p \vee \triangle \neg \triangle^{*} p \tag{4.2}$$

Unlike in the case of (4.1) where it was easy to reject the principles that made up that disjunction it does not look so easy in this case.

(4.2) follows from excluded middle and the following two claims:

$$\triangle^{*} p \rightarrow \triangle \triangle^{*} p \tag{4.3}$$

$$\neg \triangle^{*} p \rightarrow \triangle \neg \triangle^{*} p \tag{4.4}$$

Note that this is neither **4** or **5** for $\triangle$ nor **4** or **5** for $\triangle^{*}$.

4 for $\triangle^{*}$ is: ($\triangle^{*} p \rightarrow \triangle^{*} \triangle^{*} p$) which does hold. While it seems natural to allow the admitting relation for determinacy to be non-transitive Williamson shows transitivity follows from our definition $\triangle^{*}$. When we consider the infinite conjunction of *A* and determinately *A* and determinately determinately *A* ... in terms of formal semantics we introduce the admitting relation for determinately*: admits*. Determinately* α is true at a point *s* if and only if α is true at every point that *s* admits*.

> ... let a point *s* admit* a point t if and only if either *s* admits *t*, or *s* admits a point that admits *t*, or *s* admits a point that admits a point that admits *t*, or ... (Williamson, 1994, p.160)

So even though it made sense that for the supervaluationist framework the admitting relation for determinately was not transitive, given the definition of determinately* it follows that the admitting* relation is transitive and hence the **4** axiom holds for determinately*. It follows that we have "$\Delta^* \Delta^* p$ always coinciding with $\Delta^* p$, and $\Delta^* p$ will then seem to be a sharp notion" (Keefe, 2000, p.210). Williamson (1994) suggests a move that could be made here, we could express the vagueness of Δ^* in terms of a new notion 'definitely!' (see footnote 10). Keefe suggests this approach saying that we can even make this move earlier on. That is, we can express the vagueness of Δ in terms of a new operator.[14] This response commits us to a hierarchy of determinacy operators which, as I have argued above seems undesirable.

Moreover, we can avoid this need to go into hierarchies of determinacy operators if we reject one of (4.3) or (4.4). It seems that (4.3) follows from the definition of Δ^* and some plausible infinitary logic.[15] Therefore if we want to deny (4.2) this does not look like the way to deny it.[16]

If we have (4.3) and **B** then we can show that (4.4) follows by the following argument.

1. $\neg \Delta \Delta^* p \rightarrow \neg \Delta^* p$ (contrapositive of (4.3))
2. $\Delta(\neg \Delta \Delta^* p \rightarrow \neg \Delta^* p)$ by RN and 1.
3. $\Delta\neg \Delta \Delta^* p \rightarrow \Delta\neg \Delta^* p$ by K and 2.

[14] See (Keefe, 2000, pp.210-211).

[15] See Appendix B.

[16] Field (2008) denies it by giving a non-classical theory in which a conjunction of determinate truths need not be determinate, that is, he denies (4.3). However, Bacon (2020) argues that one still has to reject **B**.

4. $\neg(\triangle^* p) \rightarrow \triangle\neg\triangle(\triangle^* p)$ an instance of **B**.

5. $\neg\triangle^* p \rightarrow \triangle\neg\triangle^* p$ from 4. and 3. and transitivity

So, it looks like we are committed to (4.3) and (4.4), if **B** is, as Williamson says, plausible. Since there are good reasons to think that (4.3) holds for $\triangle^*$ it follows that if we want to deny (4.4) then it looks like the most natural way to do so is to deny **B**. As noted earlier, Williamson's frame conditions say that the admissibility relation is symmetrical and from this **B** follows. However, as Mahtani (2008), Dorr (2015) and Bacon (2020) argue, it looks like there are in fact grounds for denying **B**.

In order to look at the plausibility of these rules applying to the determinacy operator, we must take a step back and consider the vagueness framework and the accessibility relation, and Williamson's view that it is symmetric. If a point *u* is close enough to a point *v* where close enough is being less than a certain distance, then *v* is also close enough to *u*. Having said this, Williamson (1999) himself notes that these considerations in favour of **B** are not decisive.

> **B** is the least obvious of the principles listed above ... Intuitively, the crucial feature of a sorites series is that successive members differ at most slightly, and differing at most slightly seems to be a reflexive, symmetric, non-transitive relation. Although these considerations in favour of **B** are by no means decisive, we will provisionally include it as contributing to a particularly simple conception of the semantics, and call it into question if it proves to have dubious consequences. (Williamson, 1999, p.130)

The fact that **B** looks to be the problematic rule in the above argument suggests that **B** does prove to have dubious consequences. This suggests grounds for denying **B**. It would be more satisfying, however, if we had an explanation as to why it fails.

As Mahtani (2008) argues, the term 'determinately' is itself vague. Just as with other vague terms we say there are permissible precisifications but it is indeterminate which precisification we take, with the term 'determinately' it is indeterminate how we interpret it. Given the vagueness of the term 'determinately', it follows that it will be admitted by some sharpenings and not others. Mahtani argues that even if there are some languages where the meaning of determinately does not vary, these are not the languages we are interested in when considering a vagueness framework. A vagueness framework should be able to accommodate all vague sentences and we can give examples of vague sentences that contain the word determinately such as 'Determinately, Jack is bald' (p.505). The interpretations of the terms 'bald' and 'determinately' can vary and so a vagueness framework should accommodate different admissible interpretations of 'determinately' as well. Moreover, we can see that even if a sentence does not contain the word determinately, we still need to have a framework that accommodates variability of the interpretation of determinately due to higher-order vagueness.

> ... the sentence 'Jack is bald' is 2nd order precise only if the sentence'Determinately, Jack is bald' is 1st order precise. And whether 'Determinately, Jack is bald' is 1st order precise depends on whether it is true under a variety of alter-

native interpretations—including both alternative interpretations of 'bald' and alternative interpretations of 'determinately'.[17] (Mahtani, 2008, p.505)

So we can see that understanding higher-order vagueness (and higher-order precision) we need a vagueness framework that varies the interpretation of the term 'determinately'.

Putting this in the context of Williamson's distance measure, which is how he argues for symmetry, we can see that if we accept that 'determinately' can vary across interpretations then some sharpening *s* may deem another sharpening *t* admissible if *t* differs by some amount *u* but *t* may admit a wider difference. That is to say, the range of sharpenings any given point deems admissible, need not be a fixed amount. If the amount they can differ need not be a fixed amount, then the relation is not guaranteed to be symmetric, and therefore **B** need not hold as a rule. This argument seems to undermine Williamson's argument for the accessibility relation being symmetric.

If we can deny **B** then we can deny (4.4) $\neg\Delta^* p \rightarrow \Delta\neg\Delta^* p$ and thus show that we cannot derive the disjunction (4.2). This therefore blocks the argument that higher-order vagueness is simply a surface phenomena. So, for higher-order vagueness in general to make sense and for my proposed solutions for the Revenge Probabilistic Liar to work we need a logic of determinacy that rejects **B**.[18]

[17] Word changed from 'definitely' in the quote to 'determinately' to fit with my use of terminology'.

[18] A worry is that even if we can deny **B** we still have (4.3). Indeterminacy in whether it's super-determinate that $C\rho(\pi^*) < 0.5$ would require both that $\neg\Delta\Delta^*$, and that $\neg\nabla\neg\nabla^*$. But from the contrapositive of (4.3), we would get that $\neg\Delta^*\ C\rho(\pi^*) < 0.5$

4.7 Indeterminacy

In this chapter I have defended a particular type of attitude an agent ought to take towards Revenge and in previous chapters I have argued for a distinct attitude an agent should adopt towards the Probabilistic Liar. A worry for my account so far is whether we should expect the same attitude to be recommended for both the Probabilistic Liar and Revenge for the Probabilistic Liar. There seems to be a shared phenomenon in both the Probabilistic Liar and Revenge. In both cases, I have been implicitly assuming that they are both cases of indeterminacy, a view which I will defend in the next chapter.

With both these propositions, it is clear that we cannot consistently class them as 'true' or 'false' so it is natural to suggest they fall into a third category of indeterminacy. With (α), I have argued that an agent should suspend judgment where suspended judgment is understood as the agent having a determinate attitude that is represented by imprecise credences. This provides a stable belief state for the agent. Suspended judgment cannot act as a solution to (β) however. My suggested solution is instead to remove the assumption that an agent has a determinate attitude towards (β) - instead it is indeterminate what the agent's credence is towards it. This indeterminacy is also represented by imprecise credences. However, when we represent the indeterminacy of an agent's attitude the interpretation of the set is different than in the case of suspended judgment.

following from this. One option here is to deny the infinitary logic that is used to derive (4.3).

One worry then for the account is explaining and justifying having different types of credal sets towards sentences that have the same status (and importantly the agent knows or can know the status of these propositions). In the next chapter I will consider the normative question of how being certain a proposition is indeterminate ought to affect what attitude an agent ought to adopt towards that proposition.

Chapter 5
The Cognitive Role of Indeterminacy

5.1 Introduction

In the previous chapters I have presented the Probabilistic Liar and Revenge problems for the Probabilistic Liar. I have argued that rational agents ought to adopt particular attitudes to these propositions. There also seems to be a shared phenomenon in the Probabilistic Liar and Revenge problems. This suggests the wider question of whether there is some norm that governs what attitude agents ought to adopt in cases where it looks like there is a shared phenomenon. There are a variety of norms that purport to govern what attitude an agent ought to adopt depending on chance information, evidence available to the agent, accuracy considerations and so on. Being certain of the semantic status of a proposition also looks to be important information and in the case of truth and falsity there are some clear norms that we can give that govern what attitude a rational agent ought to have.[1]

> TRUTH: If an agent is certain a proposition is true she ought to believe, or be confident, or have high credence in that proposition.
>
> FALSITY: If an agent is certain a proposition is false she ought

[1] In the context of traditional belief accounts it is also said that belief aims at truth. Wedgwood (2002) for example talks about the aim of belief as a normative claim - "a belief is correct if and only if the proposition believed is true" (p.267).

> to disbelieve, be unconfident, or have low credence in that proposition.

We can understand having a low credence as having a credence less that 0.5 and having a high credence as having a credence greater than 0.5. Note that the norms I have given are *evaluative* norms as opposed to *prescriptive* norms. Using the distinction given in Mchugh (2012), a prescriptive norm tells us what one ought to believe and can guide belief, "we are accountable to them, in the sense that violating them is liable to leave us open to blame" (p.9). In contrast, evaluative norms are concerned with what is good and bad. These norms often have implications for what one ought to believe or do but this might not be in a straightforward way. So it may be that an agent S believes that *p* and this is bad because *p* is not true, but if S has been given misleading evidence for *p* she is not blameworthy for having this belief.

One might even think we can give stronger norms than these:

> TRUTH*: If an agent is certain a proposition is true she ought to believe / have credence 1 in that proposition.

> FALSITY*: If an agent is certain a proposition is false she ought to disbelieve / have credence 0 in that proposition.

Although these are stronger than the TRUTH and FALSITY norms I gave above these stronger versions also seem reasonable. If one is certain that a proposition is true being fully confident in that proposition is a reasonable response. Similarly, if one is certain a proposition is false then being fully unconfident is reasonable. Wedgwood says something similar to this:

> If a proposition is true, then the higher one's degree of credence in the proposition, the closer the belief is to being correct; and if the proposition is not true, the lower one's degree of credence, the closer the belief is to being correct.[2] (Wedgwood, 2002, p.272)

It is unclear, however, what attitude an agent ought to have towards an *indeterminate* proposition or whether there is a norm that prescribes an attitude. There are several worries with the normative question for indeterminate propositions. One is the interaction between a norm for indeterminacy and other norms of rationality. It looks like a reasonable requirement that an ideally rational agent ought to be able to meet the demands of rationality. Caie presents this as a belief norm saying:[3]

> POSSIBILITY: It must always be possible for an antecedently rational agent to continue to meet the requirements imposed on it by rationality.[4] (Caie, 2012, p.6)

Without a good reason to reject this principle, it seems we ought to try our best to not give normative principles that conflict with each other. Given problematic examples such as examples of self-reference, this

[2] Where he uses the term correct to express a normative concept. We ought to aim for correct beliefs.

[3] Caie is concerned with looking at rational agents that are not guilty of any antecedent rational failing before considering self-reference cases and the problems they pose for norms of rationality.

[4] He also gives a similar principle in (Caie, 2013, p.536): OUGHT-CAN: It must always be possible for an agent to meet the requirements imposed by rationality.

suggests we have to reject or weaken other norms of rationality.[5] With this in mind, any answer to the normative question of what attitude an agent ought to have towards indeterminate propositions also has to be compatible with other norms of belief. It is however unclear what norms should be held fixed in discussions about what attitude a rational agent ought to adopt towards indeterminacy (what I will refer to as 'the cognitive role of indeterminacy'). As I will show in section 5.3.1 some accounts of the cognitive role of indeterminacy suggest we ought to reject PROBABILISM.

> PROBABILISM: An agent's credence functions are irrational if they are probabilistically incoherent.[6]

Given Caie's POSSIBILITY how we respond to indeterminacy also opens up wider questions such as whether we ought to accept a traditional Bayesian picture or not. In light of indeterminacy, it is reasonable that the standard Bayesian picture and thus probabilism might not remain unaffected. As Field notes, the Bayesian picture of an idealised agent

> … makes sense as a crude idealization when the agent does not recognize any potential for vagueness or indeterminacy in his own sentences, but I don't think it obvious that the recognition of the possibility of vagueness or indeterminacy should leave it unaffected. (Field, 2000, p.16)

[5] POSSIBILITY is also compatible with it being impossible to ever be perfectly rational. While this is one conclusion that we can draw from POSSIBILITY and conflicts with other norms I will use POSSIBILITY as general motivation for considering weakening rationality requirements.

[6] See for example Pettigrew (2013b).

The standard Bayesian picture also assumes agents have complete sharp credences which, as I have argued for in a previous chapter, is an unrealistically strong constraint even for an ideal agent.

There are a number of accounts that have been given that give different answers to the normative question of what attitude an agent ought to adopt in cases of indeterminacy. These accounts can be categorised into two types: monistic and pluralistic accounts. That is to say, whether one thinks there is a unique attitude that an agent always ought to have towards all indeterminate propositions or whether one thinks that there is no such unique attitude.

Which side of the debate between monism and pluralism you come down on is a substantial matter. Accepting monism means accepting that there is a norm that governs indeterminacy and gives us a unique cognitive role for indeterminacy. However, within the monist position there are a number of different attitudes argued for and thus a variety of incompatible accounts. There is also a wide-ranging variety of phenomena that are widely referred to as indeterminate. Within such a broad phenomenon there are different cognitive roles for indeterminacy recommended or argued for in the literature. An advantage of the monist position would be that we learn what attitude an agent really ought to have in all of these cases once we have established it for one. It would certainly be helpful if knowing the status of a proposition helped inform us of what attitude an agent ought to have towards it. Given the existence of norms for belief when it comes to Truth and Falsity it seems reasonable to explore whether we can give a norm for indeterminacy. However, an explanation would also have to be given for why all of these cases of indeterminacy seemingly have such different cognitive roles with our intuitive responses to these cases.

We can also see that monism and certain existing norms of rationality generates a very strong theory of the cognitive role of indeterminacy. Take the example of future contingents and how strong a conclusion we can draw if we accept monism and that this is an example of indeterminacy.[7] Consider the following example:

> Will the flipped coin currently spinning in the air land heads?

When considering the attitude an agent ought to take towards this it looks like it is governed by another norm of rationality - the Principal Principle.

> PRINCIPAL PRINCIPLE: A rational agent should conform her credences to the known chances.[8]

The Principal Principle is itself somewhat contentious and there is debate about the exact nature of a chance-credence principle (see for example Pettigrew (2013c)). However, despite there being disagreement over the exact nature of a chance-credence principle we can broadly accept that there ought to be some guidance from the known chances when a rational agent forms credences about chance events.

If one accepts that future contingents are indeterminate and the Prin-

[7] See Barnes and Cameron (2008) for an argument that the open future is a type of indeterminacy. In particular, it is an example of metaphysical indeterminacy.

[8] For a more technically precise formulation see Pettigrew (2013a).

cipal Principle, then this tells us that we ought to have sharp credences to any indeterminate proposition. This alone does not pin down a unique cognitive role for indeterminacy. It entails that you should have credence c, $0 < c < 1$, but, depending on the evidence you could have credence 0.5 or credence 0.6 etc. If one is also committed to monism then it seems we have to go further and restrict the rational attitude an agent can adopt towards future contingents to some particular sharp credence. However, if we consider how I proposed we treat the Probabilistic Liar in Chapter 1 we can see that adopting such a stance towards the cognitive role of indeterminacy would prove deeply problematic if we also accept that the Probabilistic Liar is indeterminate. Pluralism offers an alternative and seems like the lesson we should take from looking at cases of indeterminacy. By looking at the insights we can get from other examples of indeterminacy and considering what the cognitive role of indeterminacy ought to be for each example of indeterminacy, we can extrapolate that the cognitive role of indeterminacy is pluralistic.

Moreover, I require a pluralist stance towards the cognitive role of indeterminacy for my arguments in previous chapters with regards to the cognitive role I suggest for the Probabilistic Liar and a distinct attitude to the Revenge problem for the Probabilistic Liar. In section 5.2 I will outline how I am committed to pluralism about the cognitive role of indeterminacy given my arguments for the Probabilistic Liar and Revenge.

In section 5.3.1 I will outline a range of accounts that have been given in the literature for the cognitive role of indeterminacy. Despite disagreeing with the methodological position of the monist, examining them highlights an important underlying point about the other normative or methodological commitments one has to have to hold particular views about the cognitive role of indeterminacy. These include important considerations like whether one wants to defend that the norm of Probabilism is universally held or ought to be dropped in cases of suspected indeterminacy or whether we should expect it always to be possible for an ideally rational agent to meet the demands of rationality. In section 5.4 I contrast the monist approach with a pluralist stance. In particular, I will draw on Williams (2012) argument that in the case of indeterminacy there is normative silence where he points to the wide range of examples of indeterminacy in the literature. Williams offers a strong version of pluralism and in section 5.4.3 I will show that this strong (or radical) pluralism faces a problem of how to account for the interaction of different norms for indeterminacy.

In section 5.5 I offer an alternative pluralist account of the cognitive role of indeterminacy - what I call Modest Pluralism. This position is one where there is no unique cognitive role for indeterminacy but there are constraints on what sorts of attitudes an agent can adopt towards indeterminacy. In section 5.6 I will revisit my commitment to pluralism in light of potential objections to the pluralist position and consider some ways my account might be amended to resolve these worries.

5.2 Probabilistic Liar, Revenge and Indeterminacy

In the previous chapters I have argued for solutions to the Probabilistic Liar and the Revenge problem for the Probabilistic Liar. In this section I will argue that both the Probabilistic Liar and Revenge are cases of indeterminacy. In section 5.2.1 I will outline what is meant by indeterminacy. In light of this understanding of indeterminacy I will revisit the Probabilistic Liar and Revenge problems in sections 5.2.2 and section 5.2.3.

I am therefore committed to the view that the Probabilistic Liar and Revenge are both cases of indeterminacy and that agents ought to adopt different attitudes towards the Probabilistic Liar and Revenge. A consequence of my view is that I must be committed to the pluralist response to the cognitive role of indeterminacy. Therefore, arguments against pluralism will act as arguments against my position.

5.2.1 Indeterminacy

There are a number of phenomena that are often referred to as indeterminate. These cases seem so wide ranging that it is unclear what the term indeterminate is picking out. It also seems like there are different categories of indeterminacy such as metaphysical indeterminacy, epistemic indeterminacy and semantic (or referential) indeterminacy. Field (1973) gives the example of referential indeterminacy when Newtonian's were talking about 'mass'. Einstein showed there are two things Newtonian 'mass' can refer to: proper mass and relativistic mass. Field argues that it " doesn't make sense to ask what physical quantity Newton and other pre-relativity physicists referred to when they used the term 'mass' " (Field, 1973, p.465). Rather he says the use of the term 'mass' was referentially indeterminate.[9] There are also examples of indeterminacy that might best be described as metaphysical indeterminacy such as quantum phenomena or phenomena involving the open future. We might think that indeterminacy is being used as an umbrella term for these different categories of indeterminacy. This might suggest that in discussions about the cognitive role of indeterminacy we ought to focus on the cognitive role of metaphysical indeterminacy and the cognitive role of semantic indeterminacy.

This approach seems unsatisfactory for a number of reasons. One thing that seems unclear on this approach is how one's identification of a phenomenon as a particular type of indeterminacy and one's ascription of the cognitive role of indeterminacy interact. Firstly, it is unclear whether there is a uniform kind of indeterminacy within each of these categories. Even within semantic indeterminacy there are a range of examples.[10] This suggests that there may be multiple cogni-

[9] Field (2000) also notes that the term 'heavier than' for pre-Newtonians is a case of referential indeterminacy. It is unclear whether this refers to the relation of having greater mass than or having greater weight than. It seems wrong to say it doesn't refer to either mass or weight, rather it is indeterminate which of these relations 'heavier than' refers to for pre-Newtonians. The referential indeterminacy for both pre-Newtonians and Newtonians points to the possibility that "future physicists may find distinctions that we miss, giving rise to indeterminacy that we can't yet be aware of" (Field, 2000, p.2).

[10] Eklund (2013) gives a number of examples of semantic indeterminacy that all seem slightly different. One example he gives is Field's example that I gave above where there exist two possible referents for the word mass. It is indeterminate whether the word mass refers to relativistic mass or proper mass. Another is an example of semantic indecision. He gives an example of a partially defined predicate 'nice'. We stipulate that "n is nice if $n < 13$, n is nice if $n > 15$". We might think this points to the meaning of 'nice' being incomplete.

tive roles of indeterminacy associated with one type of indeterminacy.

Secondly, there is potentially disagreement about what type of indeterminacy is present in certain domains, but agreement that there is indeterminacy. If we consider vagueness, there is clear disagreement between different theories of vagueness of what kind of indeterminacy it is. Some might think that there is metaphysical indeterminacy, others might think it is semantic indeterminacy and others epistemic indeterminacy but there is still general agreement that there is *indeterminacy*.

Consider the following:[11]

(†) Kilimanjaro contains more than *n* molecules.

Where we let *n* be such that it is borderline. To claim that this is a case of epistemic indeterminacy is to claim that it is either true or false that Kilimanjaro contains more than *n* molecules, but we are ignorant of (†)'s truth value.[12] The claim that this is a case of semantic indeterminacy is to say that the indeterminacy stems from some part of the expression, for example, it might be the case that the word 'Kilimanjaro' fails to pick out a unique set of molecules.[13] To say (†) is metaphysically indeterminate is to claim that Kilimanjaro is itself somehow

[11] Example from (Taylor and Burgess, 2015, p.298).

[12] See Williamson (1994) for notable version of this view. Schiffer (1999) presents a generic epistemic theory of vagueness.

[13] Supervaluationism which I discussed in the previous chapter is an account of semantic indeterminacy (see K. Fine (1975) and Keefe (2000)). For further discussion on what is meant by semantic indeterminacy see Taylor and Burgess (2015) and Kölbel (2010).

indeterminate.[14]

This leaves open the possibility that different accounts could recommend the same cognitive role towards vagueness but think this cognitive role appropriate for different reasons. One option here would be to push back and insist that if a particular cognitive role is associated with, say, metaphysical indeterminacy, then in virtue of ascribing that cognitive role to vagueness one can see that this shows us the type of indeterminacy that is *really* occurring in cases of vagueness.

If there is a unique cognitive role of metaphysical indeterminacy and the attitude is only taken towards cases of metaphysical indeterminacy this seems like a strong argument. However, as mentioned above there may be multiple cognitive roles of indeterminacy associated with metaphysical indeterminacy. As I will show in section 5.4.1 there are a wide range of examples of indeterminacy that all seem to have different cognitive roles associated with them. A number of these are examples that it would be natural to class as metaphysical indeterminacy.

Moreover, in the case of vagueness, the fact that different accounts can disagree on details of where the indeterminacy comes from, but still agree that there is indeterminacy points to the fact that there is some important agreement on the phenomena in cases of vague-

[14] One might think there are certain metaphysical facts and metaphysical indeterminacy is about such facts or one might claim that metaphysical indeterminacy does not reduce to anything but is a fundamental aspect of reality. See Barnes and Williams (2009) and Barnes (2010) for discussion of metaphysical vagueness. Barnes and Williams (2011) and Eklund (2011) for discussion of metaphysical indeterminacy.

ness. In order to consider indeterminacy in different types we would need to explain why there seems to be a shared phenomenon across different examples that people have argued are different types of indeterminacy. As well as there being debates about what type of indeterminacy is present in certain domains, there is also a debate about whether there even is metaphysical indeterminacy at all. Taylor and Burgess (2015) note that one view is that "indeterminacy always just amounts to our thought and talk failing to line up in the right way with an otherwise determinate world" (p.299). That is to say, all indeterminacy is ultimately semantic.

The disagreement in the literature suggests that we need to use indeterminacy as an umbrella term that encompasses all of these types of indeterminacy. It seems that there is some phenomena that these all have in common or that overlaps these different attempts at categorisation given the disagreement about how indeterminacy should be carved up into types. In light of this understanding of indeterminacy, I will revisit the Probabilistic Liar and Revenge problems and argue that both are cases of indeterminacy.

5.2.2 Probabilistic Liar

I have been considering the following formulation of the Probabilistic Liar. An agent Alex stipulates the meaning of a sentence (α) to be

$$(\alpha)\ \ Cr\rho(\alpha) < 0.5$$

As I have argued in Chapter 1, we can see there is an intuitive problem with the Probabilistic Liar. If Alex adopts any precise credence in (α)

this makes (α) true or false and it seems that this is in conflict with our intuitions about what credence she ought to adopt when she is certain that a proposition is true or false. We can make this intuition more precise by considering norms of belief when an agent is certain that a proposition is true or false. Consider the norms of TRUTH and FALSITY that I gave in section 5.1 and we can see that there is a conflict between these norms and any precise credence Alex might adopt.

If we accept these norms, and that the agent has knowledge of her own beliefs about (α) (i.e she can introspect on her beliefs), then we can see that if she adopts any precise credence in (α) by her own account she ought to have a different credence that the one she has adopted. If she adopts any precise credence less than 0.5 then knowing this and knowing the set up of the proposition she knows that having this attitude makes (α) true, but by TRUTH this means she ought to have a credence of greater than 0.5 in (α). If on the other hand she adopts a precise credence greater than or equal to 0.5, then she knows that this attitude makes (α) false so by FALSITY she ought to have a credence of less than 0.5 in (α). It seems that it is problematic if Alex adopts any precise credence in (α) and it is problematic if the Probabilistic Liar has either the semantic status of 'true' or 'false'. In the case of the Probabilistic Liar we can see that the attitude Alex adopts towards (α) determines the semantic status of (α).

I have argued that in the case of the Probabilistic Liar an agent ought not have any precise credence towards the proposition. In particular, I argued for the attitude of suspended judgment towards the Probabilistic Liar. Importantly suspended judgment is incompatible with talking about having any precise credence. From suspended judgment in (α) and the particular case I consider when $C(\alpha) = [0, 1]$, we can infer that $\neg C(\alpha) < 0.5$. Caie's original problem however is talking about *Cr* (the credences), rather than *C* (the credal set). The question is then whether from $C(\alpha) = [0, 1]$ we can infer that $\neg Cr(\alpha) < 0.5$? If we can make this inference then we can still run Caie's original argument. However, once we consider suspended judgment it is unclear how we can make sense of talking about the agents precise credence functions. The operator *Cr* does not make sense in the context of the agent having suspended judgment. [15]

If we take suspended judgment to be an attitude an agent can adopt towards (α) we can see that Caie's original argument in terms of precise credences becomes a reductio of the supposition that the attitude towards the Probabilistic Liar is a precise attitude. The agent does not have precise credences, they have imprecise credences which represent their comparative attitudes, and in the case of suspended judgment these imprecise credences represent that the agent has a credence gap when it comes to (α). Credence gaps can be represented by imprecise credences when we take a particular interpretation of the imprecise credal set—comparativist intersectionism—which I argued for in Chapters 2 and 3. When Alex suspends judge-

[15] The intuition here is that when an agent has $C(\alpha) = [0, 1]$ we cannot assert $\neg Cr(\alpha) < 0.5$, since $C(\alpha) = [0, 1]$ represents that there are precise credence functions $Cr_i \in C(\alpha)$ such that for some $Cr_i(\alpha) < 0.5$ and for some Cr_i, $\neg Cr_i(\alpha) < 0.5$. So from $C(\alpha) = [0, 1]$ we cannot say that she has some precise credences *Cr* and $\neg Cr(\alpha) < 0.5$ nor is it the case that $\forall Cr_i \in C(\alpha), \neg Cr_i(\alpha) < 0.5$. Given this, the idea of talking about the agent's precise credences when they suspend judgment does not seems to make sense. *Cr* unrelativized does not have a clear meaning for the intersectionist (which is how I have argued we ought to interpret the credal set).

ment in (α) what is the semantic status of (α)?

When Alex suspends judgment in (α) this seems to be a case of indeterminacy. It is not the case that Alex has made (α) true by adopting a particular credence, nor is it the case that Alex has made (α) false by adopting a particular credence. What is left but indeterminacy? Given the parallels to the Liar paradox it is useful to look to theories of truth and how the Liar paradox is accommodated. As I mentioned in Chapter 1, a type of response to the Liar paradox that I favour is a three-valued approach such as Maudlin's (2004). On this account there are three semantic categories: true, false and ungrounded. When Alex suspends judgment she does not make (α) true or false and it is therefore in some third truth category, like Maudlin's 'ungrounded'.

5.2.3 Revenge

The Revenge problem poses a new problem for what attitude an agent ought to have. It generates a new problem by drawing on the solution to the Probabilistic Liar. In the Probabilistic Liar we saw that any precise credence led to a problem for the norms of belief Probablisim and Rational Introspection. The solution of suspending judgment (or adopting imprecise credences) avoided this. If we can refer to the agent's imprecise credences however, we once again encounter a problem and the Revenge problem looks more worrisome and pervasive. Whatever solution we give, a new Revenge problem can be generated. The first Revenge problem I give as:

$$(\beta) \quad C\rho(\beta) < 0.5$$

Where C is the imprecise credal set Alex has towards (β).

As with the Probabilistic Liar, in Chapter 4 I outlined that there is an intuitive problem with whatever attitude Alex takes towards (β). This intuitive problem can be seen when we consider norms for belief. If it is determinately the case that $C\rho(\beta) < 0.5$ and Alex can introspect on her beliefs, then by the set up of the proposition she knows that this makes (β) true. By TRUTH if she is certain that a proposition is true then she ought to be confident in (β) so she ought to have $C\rho(\beta) > 0.5$. If, on the other hand, it is determinately the case that $\neg C\rho(\beta) < 0.5$ and Alex can introspect on her beliefs, then by the set of the proposition she knows that this makes (β) false. By FALSITY if she is certain that a proposition is false she ought to be unconfident in (β) so she ought to have $C\rho(\beta) < 0.5$.

This assumes Alex has a determinate attitude towards (β). If it is indeterminate what attitude she has towards (β) then there is no conflict with the TRUTH and FALSITY norms. When it is indeterminate what Alex's attitude (or credal set) for (β) is and specifically, indeterminate whether it contains credence functions that give credence less than 0.5 then it follows that (β) is indeterminate. However, unlike in the Probabilistic Liar the attitude I have argued Alex ought to adopt is not suspended judgment. When an agent suspends judgement towards (β) there is still a fact of the matter about their credal set and they have a determinate attitude towards the proposition. Instead, I argue that Alex ought to have a vague or indeterminate attitude towards (β). We need to give another cognitive response to indeterminacy: indeterminacy between different credal sets. This can be represented by an agent having a set of credal sets. If there are multiple sets of credences that might represent the agent's attitude towards (β), and it is indeterminate which of these credal sets represents Alex's attitude towards (β), then there is no fact of the matter about whether $C\rho(\beta) < 0.5$ or not. As I argued in Chapter 4, we can represent this with a supervaluationist interpretation of a set of credal sets.

This leaves us with the following two claims:

(1) The Probabilistic Liar is indeterminate and an agent ought to suspend judgment towards the Probabilistic Liar (i.e have an imprecise credal set which we interpret using the intersectionist interpretation of imprecise credences).

(2) The Revenge problem is also indeterminate and an agent ought to have a set of credal sets towards Revenge (i.e have a set of imprecise credal sets which we interpret using the supervaluationist interpretation of imprecise credences). It is indeterminate which of these credal sets represents the agent's attitude.

If we accept a pluralist approach to the cognitive role of indeterminacy then there is no issue with there being different attitudes recommended in each of these cases. We can simply accept that there are different attitudes to be had in cases of indeterminacy. If, however, one has substantial problems with the pluralist approach then this might be taken to show I ought to revise what attitude an agent is recommended to have in one of these cases. Another option would be to argue that they are not both cases of indeterminacy. In the following sections I will outline the monist and pluralist positions arguing that there are good independent reasons to adopt a pluralist stance towards the cognitive role of indeterminacy.

5.3 Monism and Pluralism

What attitude should an agent adopt towards cases of indeterminacy? In order to answer this question we have to also consider the question of whether there is a unique attitude (or type of attitude) an agent should adopt towards cases of indeterminacy. In the literature there are a variety of accounts that argue for a unique attitude that an agent ought to adopt. These monist accounts offer a range of different answers to the normative question, often with implicit assumptions in the background to support them. In section 5.3.1 I will outline some of the existing accounts of the cognitive role of indeterminacy. Monist accounts look intuitively appealing. If we have identified a range of examples as examples of indeterminacy then, if there is any normative answer to what attitude an agent ought to adopt it initially seems natural to think that this attitude would apply to all of the cases of indeterminacy.

However, given the wide range of monist accounts, accepting monism means we have to reject the majority of these accounts. While it may seem plausible to reject some of the responses to indeterminacy that have been given, as I will show in section 5.4, the wide range of examples of indeterminacy suggests that a unique attitude towards all examples of indeterminacy seems counter-intuitive. Indeed, when we consider the diversity of the phenomena of indeterminacy it looks like having a range of responses to the normative question is necessary in order to accommodate our intuitions about responses to indeterminacy across different domains. Pluralism offers an alternative, in particular, if one accepts a radical version of pluralism we can accommodate all of the different cognitive roles recommended by the various monist accounts. A radical pluralist says that in cases of indeterminacy there is normative silence. This is the claim that "*there is no doxastic attitude it is right or wrong to take to* p, when p is indeterminate" (Williams, 2012, p.222). Then, any of the norms (and associated cognitive roles) suggested for indeterminacy could be legitimate responses to the normative question.

This radical version of pluralism looks problematic however. If one accepts radically different cognitive roles of indeterminacy for two different propositions we face a new problem when considering what attitude an agent ought to take towards the conjunction or disjunction of these propositions. As I will discuss in section 5.3.1, different monist accounts have different underlying assumptions, especially about how other norms of rationality interact with the cognitive role of indeterminacy. This means that across the different accounts we see different stances being taken on the norm of Probabilism. Accepting all of these accounts therefore poses a problem for how norms of rationality interact and whether it is possible for a rational agent to meet the demands of rationality. I will argue that radical pluralism about the cognitive role would require a logical pluralism, and that this leads to problems that I call "interaction effects" between different domains.

In section 5.4 I will present Williams (2012) argument for normative silence and then show how the problem of interaction effects between mixed domain propositions arises for his account. I argue that this is a problem for radical versions of pluralism, but rather than this suggesting we adopt a monist approach it suggests we ought to adopt a modest version of pluralism. In section 5.5 I will argue for my modest version of pluralism and how it avoids the problem of interaction effects while still not pinning down a unique answer to the cognitive role of indeterminacy.

5.3.1 Monism

In this section I will outline some of the different positions in the literature to show the variety of different accounts that have been given. Accounts that offer a unique answer to the cognitive role of indeterminacy have been given by Barnett (2009), Caie (2012), Field (2000, 2003), Schiffer (2000, 2003) and Smith (2009).

If one wanted to defend the view that there is a unique cognitive role of indeterminacy, then one not only has to reject the majority of these accounts, but also explain why there is such a lack consensus between different accounts that have been offered. Caie (2012) gives the following response to indeterminacy

> INDETERMINACY: For any proposition ϕ, it is a consequence of the claim that one ought to believe that ϕ is indeterminate that one ought to be such that it is indeterminate whether one believes ϕ. (Caie, 2012, p.26)

He argues that the indeterminacy in the objects of doxastic states filters up to the doxastic states, and therefore we have indeterminate doxastic states.

Caie offers his account as an alternative to a stance he calls REJECTION. This is a position that has been given as an attitude that an agent ought to adopt towards the Liar paradox (or semantic paradoxes more generally).

> REJECTION: For any proposition ϕ, it is a consequence of the claim that one ought to believe that ϕ is indeterminate, that one ought to reject both ϕ and its negation. (Caie, 2012, p.3)

Where rejection of ϕ can be thought of as having an appropriately low credence in ϕ.[16] Field also expresses this view towards the Liar paradox saying that if we consider cases where we are certain of indeterminacy that there is support for the view that "both they and their negations are believed to degree 0" (Field, 2003, p.466).[17] Field takes this approach as a way of unifying our responses to vagueness and the semantic paradoxes. Both the semantic paradoxes and vagueness are taken as examples of indeterminacy and in both cases REJECTION is the recommended attitude.

There are also accounts that focus on what attitude an agent ought to take towards paradigmatic cases of vagueness. Barnett (2009) argues that in cases of vagueness such as the proposition 'Harry is bald'

[16] This view appears in Parsons (1984) (in response to the Liar paradox).

[17] Field also gives a weaker version of this for cases of suspected indeterminacy.

> For an agent to treat A as potentially indeterminate is for him to have degrees of belief in it and its negation that add to less than 1. (Field, 2000, p.18)

(where Harry is a borderline case of bald) it should be vague whether you believe 'Harry is bald' is true and vague whether you believe 'Harry is bald' is false. Vagueness is not ignorance entailing on this view, but rather captures that you "simply do not clearly know what is going on, in the sense that you do not clearly know that it is true that Harry is bald and you do not clearly know that it is false that Harry is bald" (Barnett, 2009, p.24).

Schiffer also focuses on the example of vagueness to give his account. He argues that an agent ought to adopt a special type of non-probabilistic partial belief for cases of vagueness. He calls this vagueness related partial belief (and refers to it as v-believing). He distinguishes vagueness related partial belief from uncertainty noting that it is not a measure of uncertainty.

> When one is confronted with what one takes to be a paradigm borderline case of a bald man, one doesn't take oneself to be uncertain as to whether or not the man is bald; that's resolved by one's taking him to be a borderline case of a bald man ...to take someone to be a borderline case of a bald man is, roughly speaking, just to v-believe that the person is a bald man. (Schiffer, 2000, p.223)

He also notes that in cases of indeterminacy generally (not just indeterminacy in cases of vagueness) an agent ought to have vagueness-related partial belief (see (Schiffer, 2003, p.178)).

Smith agrees with Schiffer's view that there is something different about the attitude one takes towards vagueness and towards uncertainty. However, unlike Schiffer he argues that we should have degrees of belief towards both cases of vagueness and uncertainty.

> … what we have is one univocal notion of degree of belief—one single system of assignments of degrees of belief to propositions—but where the degrees assigned sometimes behave in accordance with the laws of probability, and sometimes do not. (Smith, 2009, p.495)

He argues for this by giving an account of degrees of belief as expected truth value.

In these various accounts we can see there is not only a wide range of attitudes recommended but also that there is disagreement on whether an agent's attitude should be probabilistic in cases of indeterminacy. Schiffer's view clearly rejects the view that an agent should have subjective probabilities towards cases of indeterminacy. Positions such as Smith's and Field's say agents ought to have degrees of belief in cases of indeterminacy, but these degrees of belief do not obey the norm of Probabilism. There are a number of arguments for why agents ought to have probabilistically coherent belief (as I mentioned in Chapter 1). Rejecting Probabilism therefore seems like a rather strong move to make and those that are committed to the norm of Probabilism reject approaches such as REJECTION on the grounds that it violates Probabilism.

As Bacon (2018) notes, if one is committed to classical probability (and therefore Probabilism) then a natural attitude to adopt towards indeterminate propositions is to have the same attitude as when uncertain. Within a classical probability framework we can represent uncertainty by having a middling credence.

> That is, one should have the same credal attitude concerning Harry's baldness as one should have about the outcome of a coin flip, for example. For according to the probability calculus, the only alternative to having middling degrees of belief about a borderline proposition is to either assign a credence of 1 to that proposition, or to its negation——and this seems to be absurd … (Bacon, 2018, p.124).

As I have noted in previous chapters, using a middling credence to represent uncertainty is not very satisfactory. This suggests that in cases of indeterminacy there may be good reasons to reject Probabilism. If one accepts the view that agents need not have probabilistically coherent beliefs in cases of indeterminacy there is then the question of how exactly an account diverges from the norm of Probabilism. There are various ways of weakening Probabilism. I have, for example, argued for a position that rejects Probabilism, arguing that the completeness requirement on credences is too strong, even for an ideal agent. The resulting formal structures represents an agent's attitudes by credal sets which are not themselves probabilistic *but* each credence function in the credal set does obey Probabilism. While the imprecise credences reject Probabilism they stay close to it, the departure rejectionism requires is much more radical.

Another thing to note is that some of these accounts seem specifically directed at an agent's attitude towards vagueness rather than indeterminacy in general. Vagueness is often taken as a paradigmatic example of indeterminacy and if one is committed to there being a unique cognitive role of indeterminacy, then whatever attitude one ought to take towards vagueness seems like it ought to apply to indeterminacy in general. Focusing on particular cases of indeterminacy such as vagueness also shows the variety of attitudes that have been recommended even for a specific domain of indeterminacy.

5.4 Pluralism

In contrast, the other methodological approach that can be taken towards the cognitive role of indeterminacy is to embrace pluralism. On this approach we do not claim that there is one unique attitude an ideally rational agent ought to take towards indeterminacy, but rather there are multiple attitudes that it is permissible for an agent to adopt. This claim has both strong and weak readings. On the strong reading it says that in any given case of indeterminacy there are multiple attitudes it is permissible to adopt. The weak reading says that there are cases where the unique appropriate attitude to adopt is one thing, and other cases where the unique appropriate attitude is another.

Pluralism as a general approach encompasses a number of different approaches that could be taken towards the cognitive role of indeterminacy. In this section I will outline a radical version of pluralism towards the cognitive role of indeterminacy given by Williams (2012). His approach calls for complete normative silence when it comes to indeterminacy. In section 5.4.3 I will argue that this form of pluralism faces problems when we consider how different attitudes one might adopt towards indeterminacy interact with each other. I will argue that this suggests there ought to be some normative constraints on the cognitive role of indeterminacy.

As we can see there is significant disagreement between the various monist approaches, this may be explained in part by the range of phenomena that are called indeterminate. In section 5.4.1 I will focus on the variety of examples of domains of indeterminacy. These suggest that indeterminacy is a phenomenon that calls for a number of different attitudes.

5.4.1 Examples of Indeterminacy

As I argued in section 5.2.1, it seems like there may be different types of indeterminacy but that we ought to group these under an umbrella term. In this section I turn to looking at some examples of different domains of indeterminacy where the variety of cases has a variety of recommendations regarding what attitude an agent ought to have.

Williams (2012) presents a variety of cases that he notes many have categorized as cases of indeterminacy. A selection of these are as follows:

(1) Will the flipped coin currently spinning in the air land heads?

(2) Is the Liar sentence true?

(4) Is Patchy red?

(7) Is the King of France bald?

(8) Is *The Turn of the Screw* a ghost story?

(10) Is this superposed particle spin-up?

In (1) it seems appropriate to adopt a credence of 0.5 (since the known chance of a fair coin landing heads is 0.5). In the case of (2) Williams says: "The Liar sentence, at face value, entails a contradiction – so it seems we should have no more confidence in it than in contradictions". (4) is a paradigmatic borderline case so it may seem like the appropriate attitude to have is " some sort of undecided, balanced state " and (10) "To ensure the empirical adequacy of quantum physics . . . in cases of (seeming) indeterminate property attribution . . . take an intermediate value." (pp.218-9). It seems then that these intuitions are mostly fulfilled by different credal states and thus there are different cognitive roles of indeterminacy.

Given the wide range of examples of indeterminacy and the variety of attitudes it seems appropriate to adopt towards them, Williams suggests that in the case of indeterminacy there is normative silence. Normative silence says that in the case of indeterminacy there is no general norm that governs what attitude an agent ought to adopt just in virtue of being certain that the proposition is indeterminate.

> ... *so far as general alethic norms go* there are simply no constraints on what the Godlike attitude to p should be, when p is indeterminate. (Williams, 2012, p.223)

There are several ways of interpreting the normative silence claim. A strong version of normative silence says not only is there no unique norm that governs what attitude an agent ought to adopt in cases of indeterminacy, but also that *any* cognitive role of indeterminacy can be adopted. This leaves open the possibility that in some circumstances (or towards certain propositions) it would be permissible for an agent to adopt REJECTION, in other circumstances one might adopt a vagueness related partial belief, in others INDETERMINACY and so on.

Complete normative silence is one option, if there are no normative constraints on an agent's attitude towards indeterminacy then it would be permissible for an agent to take different attitudes towards two vague propositions such as 'Harry is bald' and 'Patchy is red' (where patchy is a borderline case of red and Harry a borderline case of baldness). This seems like a very radical version of the pluralist position and a more plausible version of radical pluralism is if we have a domain specific normative pluralism. The problem I will present in section 5.4.3 for domain specific normative pluralism does however also apply to this more radical pluralism so nothing hangs on rejecting this view.

A domain specific pluralism would leave open the possibility of many different norms that govern indeterminacy across the range of examples. We can therefore accept that there may be a range of cognitive roles associated with indeterminacy but that there are norms that recommend what attitude an agent ought to adopt based on the domain. We might associate a certain type of attitude with indeterminacy that arises from vagueness, another type of attitude with indeterminacy that arises from semantic paradox, another with indeterminacy from future contingents etc. This is in line with the position Williams gives. He argues that even though there is no *general* norm that constrains an agent's attitude towards indeterminate propositions, there are local norms. This draws on Maudlin's view that there are local norms that say one ought to believe the Liar sentence.

> ... even if indeterminacy is a 'normative gap' so far as alethic norms go, other more local considerations can kick in, so that all things considered one should believe the indeterminate proposition. [...] The 'local alethic norms' should be thought of as having the status of conventions. They are regularities in our attitudinal reactions to indeterminate cases, regularities that we play a role in sustaining, and that we have an interest in maintaining. There's no need to have the same convention for all cases of indeterminacy. (Williams, 2012, p.223)

In the following sections I will present a more detailed look at domain specific normative pluralism and problems for this view. In the next section I give a more detailed explanation of what this account entails.

5.4.2 Domain Specific Pluralism

While domain specific pluralism about the cognitive role has some appeal when looking at the range of examples presented in section 5.4.1 there are a number of questions about how exactly this account would work. One question is how do we know what domain a proposition is in. Another is to question why we ought to commit to pluralism if there are different norms for each domain. If we have well defined domains and can give different norms for vagueness and semantic paradox and future contingents then it seems that the umbrella term of indeterminacy is not playing a useful role in the normative question and instead we ought to think of each of these domains as different subcategories of indeterminacy. In this section I will outline how we can understand what a domain is and why re-categorising the examples above as either not cases of indeterminacy or as different types of indeterminacy is not a fruitful strategy.

The question of what makes a proposition a member of a particular domain is addressed by Lynch (2009) when considering domain specific logical pluralism. He gives a simple answer: the subject matter. Expanding on this we can say that it depends on the concepts involved in the proposition. Division of different propositions seems very natural. It is clear that '2+2=4' is clearly about mathematics and 'murder is wrong' is an ethical claim. If we think domains can be divided up in this way and that there are norms for each domain then a natural question is whether that account is really a pluralist account.

We might think that in the examples above that the use of the term indeterminacy to describe them was incorrect or misleading. Williams (2012) considers this and gives three main strategies the monist might use to resist pluralism:[18]

- **Misdiagnosis**: Not all of these are cases of indeterminacy- some have been misdiagnosed as indeterminacy and in all genuine cases of indeterminacy there is a consensus about what attitude an agent ought to have.
- **Revision**: We might have some intuitions about what attitude an agent ought to have in each of these cases but some of these intuitive responses in fact need revision if they don't fit with the attitude we recommend an agent have towards indeterminacy.

[18]See (Williams, 2012, §3).

- **Ambiguity**: 'Indeterminacy' is ambiguous and the examples given above are actually cases of different phenomena.

If we think that there has been misdiagnosis or ambiguity in the identification of the examples as cases of indeterminacy then we explain intuitions about domain specific norms for these different examples without any need to suggest pluralism.

The misdiagnosis strategy requires looking to the literature in each of the areas and establishing whether the diagnosis of classing these propositions as indeterminate is incorrect in certain situations. As Williams notes, those supportive of this strategy "will need to fight things out with the first-order literature on the future, on counterfactuals and credences, on ordinary material things, and so on" (Williams, 2012, p.222). It is certainly plausible that by using this strategy one could form the view that some of the above examples are not in fact cases of indeterminacy. The misdiagnosis method could theoretically narrow the range of examples of indeterminacy, but so long as at least there is some range of examples of indeterminacy with different cognitive roles left this does not remove the argument for normative silence.

Whether the misdiagnosis strategy has much effect will depend on what one takes *indeterminacy* to be. Why is that we think that there is a connection between the phenomena occurring in cases of vagueness, future contingents, semantic paradoxes etc. A simple response to this is just that as I have suggested earlier we use the term indeterminacy minimally to refer to the semantic status of propositions that are meaningful but not true or false. On this definition we can explain why even while pointing out how wide ranging and different these examples are they are all still cases of indeterminacy. Considering the ambiguity strategy, if we have gone through the literature and determined that some of the phenomena are not in fact cases of indeterminacy it is likely that there will still be a selection of different phenomena that are classed as indeterminate. The ambiguity response could then be considered: if we use the term indeterminacy in such a broad sense we can question whether indeterminacy is just an ambiguous phenomena and there are actually different phenomena occurring in each of these examples. Looking at the ambiguity response we can see that there are a number of different views that are compatible with it. One view is to accept that there are a number of subcategories of indeterminacy related to different domains such as vagueness or the semantic paradoxes but to still group them together as disjunctive concepts under the term indeterminacy.

Another view would be to argue that in light of there being different subcategories of indeterminacy it would be better to be discussing these subcategories of indeterminacy rather than the umbrella term. On this view we consider the semantic status of each of the different propositions to actually be different: indeterminacy_1, indeterminacy_2, indeterminacy_3 etc. and despite referring to them all under one umbrella term this is misleading (and our normative discussion ought to be focused on these subcategories). This strategy ignores the similarities between the different examples of indeterminacy however. As with my discussion of indeterminacy in section 5.2.1 all of these categories have in common the general classification of *indeterminacy*. It is often the case that when we classify a proposition as a case of indeterminacy we do not specify what we think the relevant source or logic of indeterminacy is. It seems we have a generic concept of indeterminacy and an account that denies this would have to give an explanation for why it seems like there is something in common across the supposedly distinct phenomena.[19]

What looks like a more pressing issue for domain specific normative pluralism is that it looks like endorsing domain specific normative pluralism entails one must also accept domain specific logical pluralism.[20] In particular, we can see that if we adopt the radical pluralism view described above we must also accept logical pluralism since the different normative principles that are accepted differ on their underlying logical commitments.

Suppose for example one defined logical validity in terms of its conceptual role in the regulation of belief as Field (2015) does, B follows from A if and only if there is no rational credal state where you are more confident in A than in B. If we consider different cognitive roles for different domains we can see that on the rejectionist approach one cannot rationally be more confident in A than in $\triangle A$. From the definition of validity given we can then say that $\triangle A$ follows from A. This is an inference that is characteristic of certain types of logic such as global supervaluationist logic. For another domain it might be the

[19] Barnes and Williams (2011) for example argue that we have a generic notion of indeterminacy.

[20] Drawing on Pedersen (2014) where he argues that there are links between logical pluralism, alethic pluralism and metaphysical pluralism. For example having a certain view about domain specific alethic pluralism will entail domain specific logical pluralism.

case that the cognitive role of indeterminacy is not rejectionist and so we can have a cognitive role of indeterminacy where A does not entail $\triangle A$. From this it follows that there are at least two implicit logics for different domains. In the next section I outline how this poses a problem for domain specific radical pluralism about the cognitive role of indeterminacy.

5.4.3 Interaction Effects

If domain specific pluralism about cognitive roles of indeterminacy implies that there are different probability theories (or different implicit logics) for different fragments of the language (different domains) then a substantial problem for the account is what probability theory (or logic) applies to the whole.

We encounter this problem as soon as we consider the interaction between different examples of indeterminacy from different domains. One reason why the pluralist approach looks appealing over an alternative strategy of positing multiple different semantic values for different cases of 'indeterminacy' is that multiple semantic values greatly complicate an account. Not only would we need a different semantic value for each domain but also new semantic values for any mixed domain propositions. An advantage of the pluralist answer to the normative question then is that it gives an explanation for the lack of uniformity of response to indeterminacy without suggesting that it is due to there in fact being different phenomena occurring.

However, if we endorse this radical version of pluralism it seems we re-encounter this problem in the interaction between different norms for indeterminacy, since, as we have seen there are incompatible approaches taken towards other norms of rationality by various accounts. We might for example consider the examples given in Williams (2012) (2) and (4):

(2) Is the Liar sentence true?

(4) Is Patchy red?

We could answer these questions either 'yes' or 'no'. Consider the following answers to Williams' questions:

(2a) The Liar sentence is true.

(4a) Patchy is red.

We might think that different attitudes are recommended in each of (2a) and (4a). For example, in the semantic paradox case one might think that the correct attitude to adopt is REJECTION. We can show that the Liar paradox leads to contradiction so it looks like one ought not have any more confidence in the Liar paradox than we would a contradiction. So, one might think in this case that we therefore ought to reject both the Liar and it's negation. In the vagueness case one might think that the correct attitude to adopt is Caie's INDETERMINACY approach. That is to say, in borderline cases where it is indeterminate if patchy is red it ought to be indeterminate whether one believes that patchy is red.[21]

[21] Note these are just some possible examples of what attitude one might take towards these cases. Any combination of domain and different attitude one ought to adopt to an indeterminate proposition in that domain will serve to illustrate this point, nothing in this example rides on accepting that these particular attitudes be taken to these particular examples.

Lynch's classification of domain specificity is that belonging to a particular domain is a feature of atomic propositions (at least) and a proposition will belong to a particular domain in virtue of being the sort of proposition it is (such as a proposition about quantum phenomena, or about future contingents or about vague phenomena). This of course raises the question of what norms (or logic) govern non atomic sentences that mix domains. For example, we can consider a conjunction of two domains such as the following:

(⋆) The Liar sentence is true and Patchy is red.

In this case we have a set of propositions from one domain - semantic paradox which are governed by one norm (such as REJECTION) and another set of propositions from another domain - vagueness which are governed by another norm (such as INDETERMINACY). We now have a proposition that is neither in the domain of vague propositions of the domain of semantic paradox and there is no norm yet that we have that governs this proposition. Looking at (⋆) it is unclear what type of norm we could give for this since any norm would have to have REJECTION and INDETERMINACY norms as special cases. Rather than reducing the problem by accepting all of these norms we can see that it generates many more questions and much more theory is needed to give a complete picture of norms for indeterminacy.[22]

[22] This same problem could be given if we accept the more radical version of pluralism where an agent could adopt different attitudes within the same domain. If there is one norm that governs an agent's response to 'Patchy is red' and another to 'Harry is bald' then we could then ask what attitude an agent should have towards:

(∗) Patchy is red and Harry is bald.

This suggests a serious problem for pluralism. It looks like commitment to pluralism about the cognitive role of indeterminacy is also commitment to logical pluralism and all the difficulties of different domains that this brings. It looks like accepting pluralism has amplified the problem of the normative question rather than helped, which was the original claim in its favour. This therefore suggests a problem for my account. However, as I will discuss, the problem of interaction effects only exists if we take a form of radical pluralism towards the cognitive role of indeterminacy. In the next section I will present modest pluralism which still agrees with normative silence with regards to the cognitive role of indeterminacy but in a much weaker way.

One response to these conjunction cases is to always take the weaker of the two norms. Lynch (2009) considers this as a response to a similar worry for domain specific logical pluralism:

> The weakest logic in play is that which has the fewest logical truths or which sanctions the fewest valid inferences.
>
> > MODEST: where a compound proposition or inference contains propositions from distinct domains, the default governing logic is that of the compound or inference's weakest member. (p.100)

This may seem reasonable at first glance but if we take a wide enough range of cases then it seems like we eventually endorse the weakest norm for indeterminacy which we can say applies to all cases of indeterminacy (although stronger norms *may* be given for some cases of indeterminacy). On this view then, it looks like we effectively end up with monism where there is one (very weak) norm that is minimally the sort of attitude one ought to take towards indeterminate propositions. MODEST also relies on the idea that all logics can be ordered neatly from strongest to weakest, or in the normative case that all norms can be ordered neatly from strongest to weakest. It is not clear that this will be true and as such it still potentially leaves us with lack of clarity over what norm to adopt in certain cases.

For example in Classical-Probabilism the law of excluded middle is a tautology:

$$\models \phi \vee \neg\phi$$

but also determinately ϕ is not entailed from ϕ

$$\phi \not\models \triangle\phi$$

Whereas in Kleene-Probabilism the law of the excluded middle is not a tautology:

$$\not\models \phi \vee \neg\phi$$

but determinately ϕ is entailed from ϕ:

$$\phi \models \triangle\phi$$

It does not seem like one of these is stronger than the other, rather they seem incomparable.

Williams (2014b) tackles a case of interaction effects with respect to indeterminacy in personal identity and the cognitive role an agent should adopt. He considers a number of ways indeterminacy in personal identity can arise. One way is that indeterminacy in personal identity can arise from there being possible degrees of psychological connectedness. Inwagen (1990) gives the following example:

> Suppose that a person, Alpha, enters a certain infernal philosophical engine called the Cabinet. Suppose that a person later emerges from the Cabinet and we immediately name him 'Omega'. Is Alpha Omega? (Inwagen, 1990, p.243)

He goes on to ask us to suppose that the Cabinet has been set up so as to take Alpha on indeterminate adventures (where the details of what exactly constitutes an indeterminate adventure are left up to each philosopher but is taken to be an adventure where it is not definitely true or definitely false if one would survive them). In this case Williams argues that the agent should have degree supervaluational cognitive role.

Another way there can be indeterminacy in personal identity is by fission. Parfit (1984) gives a famous example of this where he imagines he gets into a teletransporter machine which malfunctions and produces two copies of him. These post malfunction copies are not psychologically related to each other.[23] In this case, Williams argues when indeterminacy arises from fission, an agent should have a subvaluational cognitive role.

If we think that there are different cognitive roles depending on how the indeterminacy arises then we can give a pluralist account of the cognitive role of indeterminacy. However, Williams notes there will be a problem when we consider how we ought to treat mixtures of these

[23] See (Parfit, 1984, Ch. 10)

two kinds of indeterminacy. Williams gives a theory of how we can accommodate the mixing of the two types of indeterminacy for personal identity he has considered. In this theory the supervaluationist and subvaluationist accounts come out as limiting cases. While this shows that it is possible to accommodate mixed domains it is non-trivial task, and it would not just be required for one or two examples of indeterminacy but for every possible combination of indeterminacy and every new cognitive role of indeterminacy.

5.5 Modest Pluralism

Rather than suggesting one ought to adopt a monist stance towards the cognitive role of indeterminacy I take the interaction effect problem to point towards a more restrictive form of pluralism. Modest pluralism retains the idea that there is no unique cognitive role of indeterminacy and there is no attitude that an agent ought to adopt towards a proposition in virtue of being certain it is indeterminate. Unlike with radical pluralism however there are restrictions on what attitudes it is permissible for an agent to adopt towards any proposition (including indeterminate propositions).

In the various examples of different monist accounts of the cognitive role of indeterminacy there are notably different approaches towards classical Probabilism. The normative silence stance towards the cognitive role of indeterminacy can be given a weaker interpretation. Rather than leaving open the possibility of *any* norm of indeterminacy being permissible we can instead view normative silence as saying that in cases of indeterminacy there is silence between which of the preexisting *probabilistically* permissible attitudes one can take towards a proposition.

Whatever attitudes an agent might take towards indeterminacy ought to be constrained by a background theory of how we represent an agent's attitudes by probabilities. In Chapter 2 I argued for Comparativism which says that what is fundamental are an agent's comparative beliefs and that these (determinate) attitudes can be represented by an (intersectionist) imprecise credal set. This account explains how it is that numerical degrees of belief can represent an agent's determinate attitudes. It also suggests a specific weakening of Probabilism. It is an unrealistically strong constraint that an agent have a complete comparative belief ordering, hence why an agent's attitudes ought to be represented by an imprecise credal set rather than precise credence functions. Rather than requiring Probabilism, we require that precise credence functions within the credal set obey Probabilism. In the special case where the credal set comprises of just one credence function we get precise credences and Probabilism.

We can also represent when an agent credences are indeterminate using a formal framework, with a supervaluationist interpretation of a set of credal sets. Here it looks like there is no additional weakening of Probabilism since we require that any precise credence function in a credal set obeys Probabilism, so in the case where an agent has indeterminate attitudes between a set of precise credence functions, each permissible precisification obeys Probabilism. In the case where an agent has indeterminate attitudes between sets of intersectionist imprecise credences any permissible precisification is an intersectionist credal set where each credence in that intersectionist credal set obeys Probabilism.

This theory of how numerical degrees of belief can be used to represent an agent's attitudes therefore constrains the types of response that we can give to the normative question of 'what attitude ought a rational agent adopt towards indeterminate propositions?' There is room for pluralism, since there are circumstances where one might have a precise credal attitude, others where one might have an intersectionist credal set and others where one might have an indeterminate attitude towards a proposition which we can represent by taking a supervaluationist set of credal sets. Importantly all these approaches require that precise credences, when they appear in a credal set, should obey Probabilism. Recall the example of open future and the Principal Principle I gave in the introduction, this example already commits us to something like this view once we drop the assumption of monism. There are a range of possible cognitive roles an agent can adopt when they are certain of indeterminacy, but these cognitive roles are constrained.

As we can see then, with this account we avoid the interaction effect problem. Since the types of attitude that an agent can adopt are constrained by the framework of either being a determinate attitude which we can represent by an intersectionist credal set, or being an indeterminate attitude which we can represent by a supervaluationist set of credal sets. There is a clear way in which each of these account deviates from Probabilism. There are probabilistic constraints within every point in a credal set, so when there's indeterminacy over sets of credal sets we know how things have to work in every point in every precisification. Moreover, we know this continues on all set nestings, so, for every candidate credal set we know how things have to work for each precisification. As I have noted there is a principled way in which Probabilism is weakened (while retaining much of the intuition behind the argument for Probabilism, since, in the case where the agent has only one credence function in their determinate credal set and there is no indeterminacy we get precise credences.)

In the next section I will revisit the examples from Williams that helped motivate pluralism and show how on modest pluralism we still get a range of attitudes to these examples. Based on my suggestions of the attitudes an agent would adopt in these cases we can then revisit a case of mixed domains and see that there is no background issue with there being a clash between the type of approach one could take.

5.5.1 Modest Pluralism Applied

Taking some of the examples given above I suggest what sort of attitude one might take towards these examples of indeterminacy.

> (1) Will the flipped coin currently spinning in the air land heads?

In this example (assuming the coin is fair) an agent ought to apply a chance credence principle and adopt the precise credence 0.5.

> (2) Is the Liar sentence true?

As briefly mentioned earlier, it might seem like the REJECTION response would be appropriate in the case of the Liar paradox, since, at face value, it looks like the Liar sentence entails a contradiction. Endorsing the REJECTION response would clearly be problematic for modest pluralism as sketched above. Does the Liar paradox therefore pose a problem for my account? I think not, there has been extensive work on theories of truth and the Liar and as such there are a variety of accounts which accommodate the Liar in different ways. Considering an account such as Maudlin's one might think the attitude to adopt is to have a credence gap (which is representable by a credal set).

(4) Is Patchy red?

This is a case of vagueness, where we take Patchy to be a borderline case of red. I think an agent ought to have an indeterminate attitude towards whether Patchy is red. This would be best represented by having a supervaluationist set of credal sets where each credal set is a permissible precisification. It is, however, indeterminate which of the credal sets represents the agent's attitude. Revisiting a conjunction example we can also see how modest pluralism deals with these. For example:

(†) The flipped coin currently spinning in the air will land heads *and* Patchy is red.

Where the agent has credence 0.5 in the coin landing heads and has an indeterminate attitude towards patchy being red. In the case of the conjunction we could apply a rule such as using the weakest logic, and the greatest weakening of Probabilism. There is a clear ordering of which attitude weakens Probabilism the most (the supervaluationist set of credal sets) and the least (having determinate credence of 0.5 - which the agent can have on a intersectionist interpretation of imprecise credences). We can therefore employ this kind of approach (which, as I showed above, radical pluralism cannot).

5.6 Probabilistic Liar, Revenge and Indeterminacy again

My account of the attitude agents ought to adopt towards the Probabilistic Liar and Revenge requires a pluralist stance. In the above section I have outlined an account of the cognitive role of indeterminacy that can accommodate this. If however one has substantial problems with Pluralism then this poses a problem for account. If one still has a commitment to monism there are some strategies that the monist can apply in light of the examples given and the range of responses suggested to these examples. In section 5.4.2 I noted three strategies that could be used: misdiagnosis, revision and ambiguity.

If one were convinced by the monist approach these responses also provide an idea of strategies one could employ to retain some of my account in earlier chapters if one is not convinced by Pluralism about the cognitive role. In particular one could either revise what attitude an agent is recommended to have in one of these cases or argue that, in light of there being different cognitive roles (and a commitment to monism) they are not both cases of indeterminacy.

In the following subsections I sketch what I think these positions would look like.

5.6.1 Same attitude

If one holds that the Probabilistic Liar and Revenge are both indeterminate and that we should hold a monist position towards the cognitive role of indeterminacy then a cognitive role that says one should have an indeterminate belief towards both of these seems like the most plausible account to me. That is, taking a credal set that we interpret using a supervaluationist interpretation. The credal set is comprised of precise credence functions which are permissible precisifications of the agent's credence towards that proposition.

This is essentially the attitude I recommend an agent take towards Revenge. In the case of Revenge I assume that the agents determinate attitudes that form the possible precisifications are intersectionist credal sets. However, the picture for Revenge essentially looks the same. The main difference is in how one would respond to the Probabilistic Liar. In the case of the Probabilistic Liar I argued that, since the agent could see that any precise credence one took towards it makes the sentence determinately true or false (and in conflict with the precise credence they have adopted towards it). Thus, I argued for an agent suspending judgment in the Probabilistic Liar.

Alternatively, one might take the line that the agent avoids these problems if the agent has a set of permissible precise credences they take towards the Probabilistic Liar but it is indeterminate which of these credences they take. Since there is no determinate fact of the matter about which precise credence represents their attitude towards the Probabilistic Liar the credal set does not make the Probabilistic Liar true or false. Indeed, this is exactly the style of response to Revenge that I have argued for. My worry with this style of solution to the Probabilistic Liar is that fails to recognise all the types of determinate attitude an agent can adopt.

5.6.2 Not both indeterminate

The second type of response that could be made is to say there are distinct attitudes recommended in the case of the Probabilistic Liar and Revenge but that this is compatible with a monist account of the cognitive role of indeterminacy because they are not both cases of indeterminacy (or perhaps different types of indeterminacy). One could argue that:

(i) Neither of them are cases of indeterminacy.

(ii) One of them is a case of indeterminacy (and that the attitude recommended in that cases is therefore the unique cognitive role of indeterminacy).

(iii) They are different types of indeterminacy (and there are unique cognitive roles for each type of indeterminacy).

In option (i) there is no issue with the attitudes I have recommended to the Probabilistic Liar and Revenge but rather with my classification of them as indeterminate. I do however think there are good reasons to class both of these as cases of indeterminacy which I have outlined in section 5.2. In option (ii) I think the most plausible argument that could be given in line with this strategy is to reevaluate my response to the Probabilistic Liar. Rather than call Probabilistic Liar 'indeterminate' we can instead think of it as having some third truth value or being ungrounded or undefined. The attitude I recommend for the

Probabilistic Liar is suspended judgment which is a determinate attitude an agent takes towards a proposition. This determinate attitude of having a credence gap towards a proposition will not be an appropriate response to as wide a range of cases of indeterminacy as the supervaluationist interpretation of a set of credal sets I recommend for Revenge.

Again this strategy seems problematic to me since it seems there are good reasons to also think the Probabilistic Liar is a case of indeterminacy. This leaves option (iii) which I think is the most fruitful strategy. Since I have been using 'indeterminate' as an umbrella term one could argue that my categorisation of the Probabilistic Liar as indeterminate is one particular type of indeterminacy and the indeterminacy in the case of Revenge is another type of indeterminacy (i.e. an ambiguity response). We can then say that suspended judgment is the cognitive role associated with the type of indeterminacy that occurs in the case of the Probabilistic Liar and having indeterminate beliefs in the correct cognitive role associated with the type of indeterminacy associated with Revenge. This response ends up being very close to my solution of modest pluralism where we can say different cognitive attitudes might be associated with different domains.

5.7 Conclusion

In previous chapters I have argued for two types of attitude an agent ought to take towards the Probabilistic Liar and Revenge problems. This raised the general question of whether there is a norm that governs what attitude one should have in these cases - a distinctive cognitive role associated with them. In this chapter I have looked at that question by looking at the cognitive role of indeterminacy in general.

I have argued that both the Probabilistic Liar and Revenge should be considered cases of indeterminacy. This seems potentially problematic, since, one might think there ought to be a unique cognitive role associated with indeterminacy. Given the attitudes I have argued for in previous chapters my account is committed to pluralism about the cognitive role of indeterminacy. Pluralism is a plausible stance to take towards the cognitive role of indeterminacy, especially in light of the wide variety of examples of indeterminacy that we can give.

There are however a number of ways we can interpret pluralism about the cognitive role of indeterminacy. One approach is to adopt a radical pluralism according to which there are no constraints on the cognitive role of indeterminacy. I show that radical pluralism about the cognitive role of indeterminacy is a problematic position. If we accept radical pluralism then we accept that there might be a wide range of different cognitive roles of indeterminacy that one might endorse for different domains or situations. Some of these attitudes require different under underlying logical commitments. This leads to a problem for radical pluralism which I refer to as interaction effects. The radical pluralist has to explain how these different attitudes interact. We can consider non-atomic propositions where the atomic components might be given incompatible attitudes.

This suggests that we should not accept radical pluralism. I argue for a position I call modest pluralism, on this view there are constraints on the types of attitudes an agent can adopt. In particular, the view I have argued for requires that for a rational agent any credence function (be it a precise credence of part of a credal set) ought to obey the norm of Probabilism. This removes the problems of interaction effects and different underlying logical commitments that radical pluralism faced while still leaving open a range of attitudes an agent can adopt towards indeterminacy.

The attitudes I have argued an agent should adopt towards the Probabilistic Liar and Revenge are compatible with the constraints on belief that are given by modest pluralism about the cognitive role of indeterminacy. I have therefore shown that recommending different attitudes in these cases is not problematic in light of them both being cases of indeterminacy.

General Conclusion

In this thesis I have given an account of how to respond to the Probabilistic Liar. This is an example of probabilistic self-reference which can arise when we consider scenarios where an agent's attitude acts as evidence for the truth or likelihood of a proposition. The focus of this thesis has been on a formal formulation of the Probabilistic Liar and problems that arise for traditional Bayesian accounts in light of it. I have argued for two types of attitude. An agent ought to suspend judgment in the Probabilistic Liar (where suspended judgment is understood as having a credence gap). They also ought to have an indeterminate attitude towards the Revenge problem for the Probabilistic Liar (which is represented by a supervaluationist imprecise set of credal sets).

The first three chapters of the thesis set up the problem the Probabilistic Liar poses and argue for a particular solution. I argued for this in two steps. First, in chapter 1, I argue that we out to take the attitude of suspended judgment towards the Probabilistic Liar, taking inspiration from gap-style solutions to the Liar Paradox. I outline how suspended judgment ought to be understood in a traditional framework and, given these features, how it can be understood in a credal framework: as represented by imprecise credences.

Second, in chapters 2 and 3 I defend a particular interpretation of imprecise credences that backs up the particular way suspended judgment should be understood. This interpretation - comparativist intersectionism, fits with my proposed understanding of suspended judgment. We treat suspended judgement as having a determinate attitude - that of having a credence gap.

The later half of the thesis then addresses two potential problems for the account I have developed. The parallels between alethic self-reference and probabilistic self-reference point to the possibility of a Revenge problem for my account. In Chapter 4 I consider whether a Revenge problem for the Probabilistic Liar can be generated. I argue a Revenge problem can be given, and that an agent ought to have an indeterminate attitude towards the Revenge problem. This indeterminacy can be represented by imprecise credences, where a set of credal sets is interpreted by means of a supervaluationist interpretation.

Arguing for two different attitudes to adopt towards the Probabilistic Liar and Revenge problems raises the question of whether the solutions I have offered are compatible with a normative theory on the cognitive role of indeterminacy. In the final chapter of the thesis I argue that the Probabilistic Liar and Revenge for the Probabilistic Liar are both cases of indeterminacy. Furthermore, I argue that there is no unique attitude an agent ought to adopt in virtue of being certain that a proposition is indeterminate. Rather, I argue for a position I call modest pluralism which says there are a range of attitudes an agent can adopt towards cases of indeterminacy, but also gives some constraints on the attitudes a rational agent can adopt. This shows there is no conflict with the normative question and the attitudes I have recommended.

This thesis has also laid the groundwork for a number of further interesting projects. I have argued for attitudes agent's ought to adopt in cases of probabilistic self-reference. I haven't proposed a model that simultaneously solves both the Probabilistic Liar (and Revenge for the Probabilistic Liar) and the Liar (and Revenge for the Liar). The question of whether we should expect there to be broadly similar solutions remains and following this the further question of whether the solution I have offered for the Probabilistic Liar is compelling if they ought to solved together. Another natural future project is to look at how the attitudes the agent adopts should impact their decision making. I initially motivated looking at probabilistic self-reference by giving examples of 'real world' scenarios (such as those given by Greaves (2013)). It looks like we can give a 'real world' example of the Probabilistic Liar such as the *Promotion** example I gave in the introduction of Chapter 1. It is worth repeating here:

> *Promotion**:
> Alice is up for promotion. Her boss, however, is a deeply insecure type: he is more likely to promote Alice if she comes across as lacking in confidence. Furthermore, Alice is useless at play-acting, so she will come across that way iff she really does have a low degree of belief that she's going to get the promotion. Specifically, she will get the promotion if her credence in getting the promotion is less than 0.5 and she will not get the promotion if her credence in getting the promotion is greater than or equal to 0.5.

In light of my arguments in this thesis it follows that Alice ought to suspend judgment about whether she will get the promotion. This gives a recommended attitude, but it is unclear how suspended judgment should impact how she ought to behave. In the decision scenario she has to behave in a certain way - she cannot be in suspense in this.[24] What decision rule is most appropriate for suspended judgment?

I also argue for agent's having indeterminate attitudes in some scenarios, in particular, in the Revenge problems. This raises the further question of how an agent should act under indeterminacy and whether we can give 'real world' scenarios involving higher-order indeterminacy.[25] Given my arguments for modest pluralism for the cognitive role of indeterminacy there is also work to be done on how different decision rules interact.

There are also interesting further questions related to sets of sets of credences such as giving an update rule for sets of sets of credences, and the question of how we should understand the psychology of an agent who has sets of sets of credences representing their attitude.

While this thesis addresses the specific problem posed by the Probabilistic Liar, in looking at this I have touched on a range of topics and developed a number of accounts that are independently interesting. I have argued for a specific representation of the suspended

---

[24]There is a body of work on decision theory for imprecise credences and a number of decision rules offered, see for example S. Bradley (2014), R. Bradley (2017), Seidenfeld (2004) and Troffaes (2007).

[25]There is work on this in Rinard (2015) (who gives an account where it is indeterminate whether an action is permissible), and Williams (2014a) (he considers how agents should behave in cases where their survival is indeterminate).

judgment attitude, given arguments against recent problems posed for comparativism, argued for a position on how to represent higher-order indeterminacy of an agent's attitudes and given a position on the cognitive role of indeterminacy.

Appendix

A: Probabilistic Liar in terms of Credal Sets

$$(\beta) \quad C\rho(\beta) < 0.5$$

Where C is the credal set Alex has and $C\rho(\beta) < 0.5$ can be understood as expressing that for every credence function $Cr_i \in C, Cr_i\rho(\beta) < 0.5$.

As in Caie's argument we can say:

$$(1) \;\; \rho(\beta) = \rho\text{`}C\rho(\beta) < 0.5\text{'}$$

We must also consider the relevant introspection principles for imprecise credences.

The positive and negative introspection principles given by Caie for precise credences are as follows:

$$(2) \;\; [Cr\rho(\alpha) < 0.5] \rightarrow [Cr(\rho\text{ `}Cr\rho(\alpha) < 0.5\text{'}) > 0.5]$$

$$(3) \;\; [\neg Cr\rho(\alpha) < 0.5] \rightarrow [Cr(\rho\text{ `}\neg Cr\rho(\alpha) < 0.5\text{'}) > 0.5]$$

When $C\rho(\beta) < 0.5$ it follows that every precise credence function in the credal set gives a credence of less than 0.5 to (β). It therefore seems we can take the positive introspection principle (2) and apply this directly to the imprecise credal set. (2) tells us that if an agent has a precise credence $Cr\rho(\alpha) < 0.5$ when she introspects she is confident that her precise credence is such that $Cr\rho(\alpha) < 0.5$. Similarly, we can say when all of the precise credence functions in the credal set say $Cr_i\rho(\beta) < 0.5$ that when she introspects she is confident of this. We can give the following version in terms of imprecise credences:

$$(2^*)\ \ [C\rho(\beta) < 0.5] \rightarrow [C(\rho\ `C\rho(\beta) < 0.5') > 0.5]$$

When $\neg C\rho(\beta) < 0.5$ it follows that it is not the case that every precise credence function in C gives a credence of less than 0.5. The negative introspection principle given in (3) tells us that if it is not the case that an agent has precise credence $Cr\rho(\alpha) < 0.5$ when she introspection she is confident that it is not the case that her precise credence is $Cr\rho(\alpha) < 0.5$. Similarly, then it looks like we can say that when it's not the case that all the precise credence functions in the credal set say $Cr_i\rho(\beta) < 0.5$ that when she introspects she is confident of this. We can give the following version in terms of imprecise credences:

$$(3^*)\ \ [\neg C\rho(\beta) < 0.5] \rightarrow [C(\rho\ `\neg C\rho(\beta) < 0.5') > 0.5]$$

Assume:

$$(4)\ \ C\rho(\beta) < 0.5$$

Then by (1) $C(\rho`C\rho(\beta) < 0.5') < 0.5$

but given our assumption of (2^*) we get a contradiction.

So it follows that

$$(5)\ \ \neg C\rho(\beta) < 0.5$$

Then by (1) we get $\neg C(\rho\ `C\rho(\beta) < 0.5') < 0.5$

and by (3*) we get that $C(\rho\ `\neg C\rho(\beta) < 0.5') > 0.5$.

i.e. of the form $\neg C(B) < 0.5$ and $C(\neg B) > 0.5$.

$\neg C(B) < 0.5$ tells us that it's not the case that all the credence functions in C give credence less than 0.5 to B. From this it follows that there is at least one credence function $Cr_i \in C$, $Cr_i(B) \geq 0.5$.

$C(\neg B) > 0.5$ tells us that all the credence functions in the credal set give a credence greater than 0.5 to $\neg B$. From this is follows that $Cr_i \in C$, $Cr_i(\neg B) > 0.5$.

Which give a contradiction since we assume all the credence functions in the credal set are probabilistically coherent.

So if we can reformulate the Probabilistic Liar in terms of credal sets and accept these versions of positive and negative rational introspection it looks problematic .

B: Infinatary Logic

It seems that

$$\triangle^* p \to \triangle \triangle^* p$$

follows from the definition of $\triangle^*$ and some plausible infinitary logic.

Bacon (2020) shows that it follows from the following principles:

C1. $\bigwedge_{i<\omega} \phi_i \to \phi_n$ for each $n < \omega$

C2. $\bigwedge_{i<\omega} (\phi \to \psi_i) \to (\phi \to \bigwedge_{i<\omega} \psi_i)$

C3. If $\vdash \phi_i$ for each $i < \omega$, $\vdash \bigwedge_{i<\omega} \phi$

D1. $\bigwedge_{i<\omega} \triangle\phi_i \to \triangle \bigwedge_{i<\omega} \phi_i$

To see the plausibility of these principles consider the following (which are instances of these principle in finite classical logic):

$C1'$ $(A_1 \wedge A_2) \to A_i$ for $A_i \in \{A_1, A_2\}$

$C2'$ $((A \to B_1) \wedge (A \to B_2)) \to (A \to (B_1 \wedge B_2))$

$C3'$ If $\vdash A_1$ and $\vdash A_2$, $\vdash A_1 \wedge A_2$

$D1'$ $(\triangle A_1 \wedge \triangle A_2) \to \triangle(A_1 \wedge A_2)$

$C1'$ is conjunction elimination, $C2'$ says that from a conjunction of conditionals you can infer the conjunction of the consequent and $C3'$ is conjunction introduction. Each of these look like plausible principles of finite logic and $C1 - C3$ are generalisations of these to the infinite case. Bacon shows that $D1$ can be derived from $C1 - C3$, Necessitation, Closure and Brouwer's Principle (see Theorem 6 (Bacon, 2020, p.28)).

Given this, $C1 - C3$ and $D1$ look very plausible for the infinite case and unless someone gives a compelling argument to reject them we should accept them for now.

C1-C3 gives us: $\vdash \bigwedge_{n<\omega} \triangle^n\phi \to \bigwedge_{n<\omega} \triangle^{n+1}\phi$.

By D1 we can then infer: $\vdash \bigwedge_{n<\omega} \triangle^n\phi \to \triangle \bigwedge_{n<\omega} \triangle^n\phi$.

i.e. $\triangle^* p \to \triangle \triangle^* p$

References

Abellan, Joaquin and Serafin Moral (2003). *Maximum of entropy for credal sets*. In: International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems 11.05, pp. 587–597.

Alon, Shiri and Ehud Lehrer (2014). *Subjective multi-prior probability: A representation of a partial likelihood relation*. In: Journal of Economic Theory 151, pp. 476–492.

Anscombe, Francis J, Robert J Aumann, et al. (1963). *A definition of subjective probability*. In: Annals of mathematical statistics 34.1, pp. 199–205.

Bacon, Andrew (2018). *Vagueness and Thought*. Oxford University Press.

– (Sept. 2020). *VIII—Vagueness at Every Order*. In: Proceedings of the Aristotelian Society 120.2, pp. 165–201. DOI: `10.1093/arisoc/aoaa011`.

Barnes, Elizabeth (2010). *Ontic Vagueness: A Guide for the Perplexed*. In: Noûs 44.4, pp. 601–627. DOI: `10.1111/j.1468-0068.2010.00762.x`.

Barnes, Elizabeth and Ross Cameron (2008). *The Open Future: Bivalence, Determinism and Ontology*. In: Philosophical Studies 146.2, pp. 291–309. DOI: `10.1007/s11098-008-9257-6`.

Barnes, Elizabeth and J. Robert G. Williams (2009). *Vague Parts and Vague Identity*. In: Pacific Philosophical Quarterly 90.2, pp. 176–187. DOI: 10.1111/j.1468-0114.2009.01335.x.

– (2011). *A Theory of Metaphysical Indeterminacy*. In: Oxford Studies in Metaphysics volume 6. Ed. by Karen Bennett and Dean W. Zimmerman. Oxford University Press, pp. 103–148.

Barnett, David (2009). *Is Vagueness Sui Generis?* In: Australasian Journal of Philosophy 87.1, pp. 5–34. DOI: 10.1080/00048400802237376.

Bergmann, Michael (2005). *Defeaters and Higher-Level Requirements*. In: Philosophical Quarterly 55.220, pp. 419–436.

Bogart, Kenneth P. (1973). *Preference structures I: Distances between transitive preference relations*. In: The Journal of Mathematical Sociology 3.1, pp. 49–67. DOI: 10.1080/0022250X.1973.9989823.

Bolander, Thomas (2017). *Self-Reference*. In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Fall 2017. Metaphysics Research Lab, Stanford University.

Bradley, Richard (2017). *Decision Theory with a Human Face*. Cambridge University Press.

Bradley, Seamus (2014). *How to choose among choice functions*. In: Episteme 11, pp. 277–289.

– (2016). *Imprecise Probabilities*. In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Winter 2016. Metaphysics Research Lab, Stanford University.

Caie, Michael (2012). *Belief and Indeterminacy*. In: Philosophical Review 121.1, pp. 1–54. DOI: 10.1215/00318108-1426355.

– (2013). *Rational Probabilistic Incoherence*. In: Philosophical Review 122.4, pp. 527–575.

Camerer, Colin and Martin Weber (1992). *Recent developments in modeling preferences: Uncertainty and ambiguity*. In: Journal of risk and uncertainty 5.4, pp. 325–370.

Campbell-Moore, Catrin (2015a). *How to Express Self-Referential Probability. A Kripkean Proposal*. In: The Review of Symbolic Logic 8.4, pp. 680–704. DOI: 10.1017/S1755020315000118.

– (2015b). *Rational Probabilistic Incoherence? A Reply to Michael Caie*. In: Philosophical Review 124.3, pp. 393–406.

Chater, Nick, Noah Goodman, Thomas L. Griffiths, Charles Kemp, Mike Oaksford, and Joshua B. Tenenbaum (2011). *The Imaginary Fundamentalists: The Unshocking Truth About Bayesian Cognitive Science*. In: Behavioral and Brain Sciences 34.4, pp. 194–196. DOI: 10.1017/s0140525x11000239.

Cook, Roy T. (2007). *Embracing Revenge: On the Indefinite Extendibility of Language*. In: Revenge of the Liar: New Essays on the Paradox. Ed. by J. C. Beall. Oxford University Press, p. 31.

Cozic, Mikaël and Brian Hill (2015). *Representation Theorems and the Semantics of Decision-Theoretic Concepts*. In: Journal of Economic Methodology 22.3, pp. 292–311. DOI: 10.1080/1350178X.2015.1071503.

de Finetti, Bruno (1992). *Foresight: Its logical laws, its subjective sources*. In: Breakthroughs in statistics. Springer, pp. 134–174.

– (2017). *Theory of Probability: A Critical Introductory Treatment*. Vol. First edition 1974. New York: John Wiley.

DiBella, Nicholas (2018). *Qualitative probability and infinitesimal probability*. In: Draft of 9.7, p. 18.

Dorr, Cian (2015). *HOW VAGUENESS COULD CUT OUT AT ANY ORDER*. In: The Review of Symbolic Logic 8.1, pp. 1–10. DOI: 10.1017/S175502031400032X.

Egan, Andy and Adam Elga (2005). *I Can't Believe I'm Stupid*. In: Philosophical Perspectives 19.1, pp. 77–93.

Eklund, Matti (2011). *Being Metaphysically Unsettled: Barnes and Williams on Metaphysical Indeterminacy and Vagueness*. In: Oxford Studies in Metaphysics 6, p. 6. DOI: 10.1093/acprof:oso/9780199603039.003.0004.

– (2013). *Metaphysical Vagueness and Metaphysical Indeterminacy*. In: Metaphysica 14.2, pp. 165–179. DOI: 10.1007/s12133-013-0119-0.

Elliott, Edward (2019). *Betting Against the Zen Monk: On Preferences and Partial Belief*. In: Synthese 198.4, pp. 3733–3758. DOI: 10.1007/s11229-019-02308-4.

– (2020a). *'Ramseyfying' Probabilistic Comparativism*. In: Philosophy of Science 87.4, pp. 727–754. DOI: 10.1086/709785.

– (2020b). *Comparativism and the Measurement of Partial Belief*. In: Erkenntnis, pp. 1–28. DOI: 10.1007/s10670-020-00329-x.

– (2020c). *What Is 'Real' in Interpersonal Comparisons of Confidence*. In: Australasian Journal of Philosophy 0.0, pp. 1–15. DOI: 10.1080/00048402.2020.1849326.

– (ms). *Extract from 'Credal Sets and Strength of Belief'*.

Ellsberg, Daniel (1961). *Risk, ambiguity, and the Savage axioms*. In: The quarterly journal of economics, pp. 643–669.

Eriksson, Lina and Alan Hájek (2007). *What Are Degrees of Belief?* In: Studia Logica 86.2, pp. 185–215.

Feldman, Richard (2000). *The Ethics of Belief*. In: Philosophy and Phenomenological Research 60.3, pp. 667–695. URL: http://www.jstor.org/stable/2653823.

Field, Hartry (1973). *Theory Change and the Indeterminacy of Reference*. In: Journal of Philosophy 70.14, pp. 462–481. DOI: 10.2307/2025110.

– (1978). *Mental Representation*. In: Erkenntnis 13.July, pp. 9–61. DOI: 10.1007/BF00160888.

– (2000). *Indeterminacy, Degree of Belief, and Excluded Middle*. In: Noûs 34.1, pp. 1–30. DOI: 10.1111/0029-4624.00200.

Field, Hartry (2003). *No Fact of the Matter*. In: Australasian Journal of Philosophy 81.4, pp. 457–480. DOI: 10.1080/713659756.

– (2008). *Saving Truth From Paradox*. Oxford University Press.

– (2015). *What Is Logical Validity?* In: Foundations of Logical Consequence. Ed. by Colin R. Caret and Ole T. Hjortland. Oxford University Press.

Fine, Kit (1975). *Vagueness, Truth and Logic*. In: Synthese 30.3-4, pp. 265–300. DOI: 10.1007/BF00485047.

Fine, Terrence (1971). *Rational decision making with comparative probability*. In: 1971 IEEE Conference on Decision and Control. IEEE, pp. 355–356.

– (1973). *Theories of probability: An examination of foundations*. Academic Press.

Fishburn, Peter C. (1986). *The Axioms of Subjective Probability*. In: Statistical Science 1.3, pp. 335–345.

Fodor, Jerry A. (1975). *The Language of Thought*. Harvard University Press.

Friedman, Jane (2013a). *Rational Agnosticism and Degrees of Belief*. In: Oxford Studies in Epistemology 4, p. 57.

– (2013b). *Suspended Judgment*. In: Philosophical Studies 162.2, pp. 165–181.

– (2017). *Why Suspend Judging?* In: Noûs 51.2, pp. 302–326.

Garson, James (2021). *Modal Logic*. In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Summer 2021. Metaphysics Research Lab, Stanford University.

Greaves, Hilary (2013). *Epistemic Decision Theory*. In: Mind 122.488, pp. 915–952.

Griffiths, Thomas L, Nick Chater, Dennis Norris, and Alexandre Pouget (2012). *How the Bayesians got their beliefs (and what those beliefs actually are): comment on Bowers and Davis (2012)*. In.

Hájek, Alan (1998). *Agnosticism Meets Bayesianism*. In: Analysis 58.3, pp. 199–206.

– (2003). *What Conditional Probability Could Not Be*. In: Synthese 137.3, pp. 273–323. DOI: 10.1023/b:synt.0000004904.91112.16.

– (2008). *Arguments For, or Against, Probabilism?* In: The British Journal for the Philosophy of Science 59.4, pp. 793–819. ISSN: 00070882, 14643537. URL: http://www.jstor.org/stable/40072312.

– (2009). *Arguments For—Or Against—Probabilism?* In: British Journal for the Philosophy of Science. Ed. by Franz Huber and Christoph Schmidt-Petri. Springer, pp. 229–251.

Hansson, Sven Ove and Till Grüne-Yanoff (2021). *Preferences*. In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Fall 2021. Metaphysics Research Lab, Stanford University.

Hartmann, Stephan and Wouter Meijs (2012). *Walter the Banker: The Conjunction Fallacy Reconsidered*. In: Synthese 184.1, pp. 73–87. DOI: 10.1007/s11229-009-9694-6.

Hawthorne, James (2016). *A Logic of Comparative Support: Qualitative Conditional Probability Relations Representable by Popper Functions*. In: The Oxford Handbook of Probability and Philosophy. Ed. by Alan Hájek and Christopher Hitchcock. Oxford: Oxford University Press. Chap. 13.

Inwagen, Peter Van (1990). *Material Beings*. Ithaca: Cornell University Press.

James, William (1956). *The Will to Believe and Other Essays in Popular Philosophy, and Human Immortality*. Dover Publications.

Joyce, James (1998). *A Nonpragmatic Vindication of Probabilism*. In: Philosophy of Science 65.4, pp. 575–603.

– (2005a). *How Probabilities Reflect Evidence*. In: Philosophical Perspectives 19.1, pp. 153–178.

– (2005b). *How Probabilities Reflect Evidence*. In: Philosophical Perspectives 19.1, pp. 153–179. DOI: 10.1111/j.1520-8583.2005.00058.x.

– (2010). *A Defense of Imprecise Credences in Inference and Decision Making*. In: Philosophical Perspectives 24.1, pp. 281–323. DOI: 10.1111/j.1520-8583.2010.00194.x.

Kahneman, Daniel, Paul Slovic, and Amos Tversky (1982). *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press.

Kaplan, Mark (1996). *Decision Theory as Philosophy*. Cambridge University Press.

– (2010). *In Defense of Modest Probabilism*. In: Synthese 176.1, pp. 41–55. DOI: 10.1007/s11229-009-9483-2.

Keefe, Rosanna (2000). *Theories of Vagueness*. Cambridge University Press.

Keynes, John Maynard (1921). *A Treatise on Probability*. Dover Publications.

Kölbel, Max (2010). *Vagueness as Semantic*. In: Cuts and Clouds: Vaguenesss, its Nature and its Logic. Ed. by Richard Dietz and Sebastiano Moruzzi. Oxford University Press.

Konek, Jason (2019). *Comparative Probabilities*. In: The Open Handbook of Formal Epistemology. Ed. by Richard Pettigrew and Jonathan Weisberg. PhilPapers Foundation. Chap. 6, pp. 267–348.

Konek, Jason and Benjamin Levinstein (2017). *The Foundations of Epistemic Decision Theory*. In: Mind, fzw044. DOI: 10.1093/mind/fzw044.

Kraft, Charles H, John W Pratt, and Abraham Seidenberg (1959). *Intuitive probability on finite sets*. In: The Annals of Mathematical Statistics 30.2, pp. 408–419.

Krantz, David, Duncan Luce, Patrick Suppes, and Amos Tversky (1971). *Foundations of Measurement, Vol. I: Additive and Polynomial Representations*. New York Academic Press.

Kripke, Saul A. (1975). *Outline of a Theory of Truth*. In: Journal of Philosophy 72.19, pp. 690–716.

Leitgeb, Hannes (2014). *The Stability Theory of Belief*. In: Philosophical Review 123.2, pp. 131–171. DOI: 10.1215/00318108-2400575.

Levi, Isaac (1974). *On Indeterminate Probabilities*. In: Journal of Philosophy 71.13, pp. 391–418.

– (1980). *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. MIT Press.

Lewis, David (1974). *Radical Interpretation*. In: Synthese 27.July-August, pp. 331–344. DOI: 10.1007/BF00484599.

List, Christian (2003). *Are Interpersonal Comparisons of Utility Indeterminate?* In: Erkenntnis 58.2, pp. 229–260. DOI: 10.1023/A:1022094826922.

Lynch, Michael P. (2009). *Truth as One and Many*. Clarendon Press.

Maher, Patrick (1993). *Betting on Theories*. Cambridge University Press.

– (2006). *Book Review: David Christensen. Putting Logic in its Place: Formal Constraints on Rational Belief*. In: Notre Dame Journal of Formal Logic 47.1, pp. 133–149. DOI: 10.1305/ndjfl/1143468316.

Mahtani, Anna (2008). *Can Vagueness Cut Out at Any Order?* In: Australasian Journal of Philosophy 86.3, pp. 499–508. DOI: 10.1080/00048400802001954.

– (2019a). *Imprecise Probabilities*. In: The Open Handbook of Formal Epistemology. Ed. by Richard Pettigrew and Jonathan Weisberg. PhilPapers Foundation, pp. 107–130.

– (2019b). *Vagueness and Imprecise Credence*. In: Vagueness and Rationality in Language Use and Cognition. Ed. by Richard Dietz. Springer Verlag, pp. 7–30.

Makinson, D. C. (1965). *The Paradox of the Preface*. In: Analysis 25.6, pp. 205–207. ISSN: 00032638, 14678284. URL: http://www.jstor.org/stable/3326519.

Maudlin, Tim (2004). *Truth and Paradox: Solving the Riddles*. Oxford University Press.

– (2007). *Reducing Revenge to Discomfort*. In: Revenge of the Liar: New Essays on the Paradox. Ed. by J. C. Beall. Oxford University Press.

Mchugh, Conor (2012). *The Truth Norm of Belief*. In: Pacific Philosophical Quarterly 93.1, pp. 8–30. DOI: 10.1111/j.1468-0114.2011.01413.x.

Meacham, Christopher J. G. and Jonathan Weisberg (2011). *Representation Theorems and the Foundations of Decision Theory*. In: Australasian Journal of Philosophy 89.4, pp. 641–663. DOI: 10.1080/00048402.2010.510529.

Monton, Bradley (1998). *Bayesian Agnosticism and Constructive Empiricism*. In: Analysis 58.3, pp. 207–212. DOI: 10.1111/1467-8284.00123.

Moss, Sarah (forthcoming). *Global Constraints on Imprecise Credences: Solving Reflection Violations, Belief Inertia, and Other Puzzles*. In: Philosophy and Phenomenological Research. DOI: 10.1111/phpr.12703.

Nehring, Klaus (2009). *Imprecise probabilistic beliefs as a context for decision-making under ambiguity*. In: Journal of Economic Theory 144.3, pp. 1054–1091.

Parfit, Derek (1984). *Reasons and Persons*. Oxford University Press.

Parsons, Terence (1984). *Assertion, Denial, and the Liar Paradox*. In: Journal of Philosophical Logic 13.2, pp. 137–152. DOI: 10.1007/BF00453019.

Pedersen, Nikolaj Jang Lee Linding (2014). *Pluralism ×3: Truth, Logic, Metaphysics*. In: Erkenntnis 79.S2, pp. 259–277. DOI: 10.1007/s10670-013-9476-x.

Pettigrew, Richard (2013a). *A New Epistemic Utility Argument for the Principal Principle*. In: Episteme 10.1, pp. 19–35. DOI: 10.1017/epi.2013.5.

– (2013b). *Epistemic Utility and Norms for Credences*. In: Philosophy Compass 8.10, pp. 897–908. DOI: 10.1111/phc3.12079.

– (2013c). *What Chance-Credence Norms Should Not Be*. In: Noûs 47.3, pp. 177–196. DOI: 10.1111/nous.12047.

– (2016). *Accuracy and the Laws of Credence*. Oxford University Press UK.

Raatikainen, Panu (2021). *Gödel's Incompleteness Theorems*. In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Spring 2021. Metaphysics Research Lab, Stanford University.

Rabinowicz, Wlodek (2017). *From Values to Probabilities*. In: Synthese 194.10, pp. 3901–3929.

Ramsey, Frank Plumpton (1926). *Truth and Probability*. In: Philosophy of Probability: Contemporary Readings. Ed. by Antony Eagle. Routledge, pp. 52–94.

– (1929). *Probability and Partial Belief*. In: Philosophical Papers. Ed. by Frank Ramsey. Cambridge University Press, pp. 95–96.

Rescorla, Michael (2019). *The Language of Thought Hypothesis*. In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Summer 2019. Metaphysics Research Lab, Stanford University.

Rinard, Susanna (2015). *A Decision Theory for Imprecise Probabilities*. In: Philosophers' Imprint 15.

– (2017). *Imprecise Probability and Higher Order Vagueness*. In: Res Philosophica 94.2, pp. 257–273.

Robbins, Lionel (1932). *An essay on the nature and significance of economic science*. Ludwig von Mises Institute.

Russell, Bertrand (1906). *On Some Difficulties in the Theory of Transfinite Numbers and Order Types*. In: Proceedings of the London Mathematical Society 4.14, pp. 29–53. DOI: 10.2307/2011035.

Savage, Leonard J. (1954). *The Foundations of Statistics*. Wiley Publications in Statistics.

Schiffer, Stephen (1999). *The Epistemic Theory of Vagueness*. In: Philosophical Perspectives 13, pp. 481–503. ISSN: 15208583, 17582245. URL: http://www.jstor.org/stable/2676114.

– (2000). *Vagueness and Partial Belief*. In: Noûs 34.s1, pp. 220–257. DOI: 10.1111/0029-4624.34.s1.22.

– (2003). *The Things We Mean*. Oxford University Press.

Scott, Dana (1964). *Measurement structures and linear inequalities*. In: Journal of mathematical psychology 1.2, pp. 233–247.

Seidenfeld, T. (2004). *A Contrast Between Two Decision Rules for Use with (Convex) Sets of Probabilities: Γ-Maximin Versus E-Admissibilty*. In: Synthese 140.1-2, pp. 69–88. DOI: 10.1023/B:SYNT.0000029942.11359.8d.

Skyrms, Brian (1977). *Resiliency, Propensities, and Causal Necessity*. In: Journal of Philosophy 74.11, pp. 704–713. DOI: 10.2307/2025774.

Smith, Nicholas J. J. (2009). *Degree of Belief is Expected Truth Value*. In: Cuts and Clouds. Vaguenesss, its Nature and its Logic. Ed. by Sebastiano Moruzzi and Richard Dietz. Oxford University Press, pp. 491–506.

Staffel, Julia (2019a). *Credences and Suspended Judgments as Transitional Attitudes*. In: Philosophical Issues 29.1, pp. 281–294. DOI: 10.1111/phis.12154.

– (2019b). *Unsettled Thoughts: A Theory of Degrees of Rationality*. Oxford University Press. DOI: 10.1093/oso/9780198833710.001.0001.

Stefánsson, H. Orri (2017). *What is "Real" in Probabilism?* In: Australasian Journal of Philosophy 95.3, pp. 573–587.

– (2018). *On the Ratio Challenge for Comparativism*. In: Australasian Journal of Philosophy 96.2, pp. 380–390. DOI: 10.1080/00048402.2017.1349159.

Sturgeon, Scott (2008). *Reason and the Grain of Belief*. In: Noûs 42.1, pp. 139–165.

– (2010). *Confidence and Coarse-Grained Attitudes*. In: Oxford Studies in Epistemology. Ed. by T. Szabo Gendler and J. Hawthorne. Oxford University Press, pp. 3–126.

Tarski, Alfred (1933). *The concept of truth in the languages of the deductive sciences*. In: Prace Towarzystwa Naukowego Warszawskiego, Wydzial III Nauk Matematyczno-Fizycznych 34.13-172.

Taylor, David E. and Alexis Burgess (2015). *What in the World Is Semantic Indeterminacy?* In: Analytic Philosophy 56.4, pp. 298–317. DOI: 10.1111/phib.12068.

Titelbaum, Michael (2019). *Precise Credences*. In: The Open Handbook of Formal Epistemology. Ed. by Richard Pettigrew and Jonathan Weisberg. PhilPaper Foundation, pp. 1–55.

Troffaes, Matthias CM (2007). *Decision making under uncertainty using imprecise probabilities*. In: International journal of approximate reasoning 45.1, pp. 17–29.

Vallinder, Aron (2018). *Imprecise Bayesianism and Global Belief Inertia*. In: British Journal for the Philosophy of Science 69.4, pp. 1205–1230. DOI: 10.1093/bjps/axx033.

van Fraassen, Bas (1990). *Figures in a Probability Landscape*. In: Truth or Consequences. Ed. by J. Dunn and A. Gupta. Kluwer Academic Publishers, pp. 345–356.

– (1998). *The Agnostic Subtly Probabilified*. In: Analysis 58.3, pp. 212–220.

Varzi, Achille C. (2007). *Supervaluationism and Its Logics*. In: Mind 116.463, pp. 633–675. ISSN: 00264423, 14602113. URL: http://www.jstor.org/stable/30163530.

Vineberg, Susan (2016). *Dutch Book Arguments*. In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Spring 2016. Metaphysics Research Lab, Stanford University.

Walley, Peter (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman & Hall.

Weatherson, Brian (2016a). *David Lewis*. In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Winter 2016. Metaphysics Research Lab, Stanford University.

– (2016b). *Games, Beliefs and Credences*. In: Philosophy and Phenomenological Research 92.2, pp. 209–236. DOI: 10.1111/phpr.12088.

Wedgwood, Ralph (2002). *The Aim of Belief*. In: Philosophical Perspectives 16, pp. 267–97. DOI: 10.1111/1468-0068.36.s16.10.

Williams, J. Robert G. (2012). *Indeterminacy and Normative Silence*. In: Analysis 72.2, pp. 217–225. DOI: 10.1093/analys/ans034.

– (2014a). *Decision-Making Under Indeterminacy*. In: Philosophers' Imprint 14.

– (2014b). *Nonclassical Minds and Indeterminate Survival*. In: Philosophical Review 123.4, pp. 379–428. DOI: 10.1215/00318108-2749710.

Williamson, Timothy (1994). *Vagueness*. Routledge.

– (1999). *On the Structure of Higher-Order Vagueness*. In: Mind 108.429, pp. 127–143. DOI: 10.1093/mind/108.429.127.

Wolford, George, Holly A Taylor, and J Robert Beck (1990). *The conjunction fallacy?* In: Memory & Cognition 18.1, pp. 47–53.

Zynda, Lyle (2000). *Representation Theorems and Realism About Degrees of Belief*. In: Philosophy of Science 67.1, pp. 45–69. DOI: 10.1086/392761.