

On Uncertainty

Brian Weatherson BA (Hons)

Department of Philosophy

Monash University

Dissertation submitted for degree of Doctor of Philosophy on 21

July 1998

Table of Contents

Synopsis	vi
Declaration of Originality	viii
Acknowledgments	ix
Chapter 1 What Probability Isn't	1
1.1 Introduction	1
1.2 The Probability Calculus	3
1.3 Probability is not Frequency	6
1.4 Probability is not Modal Frequency	10
1.5 Probability is not Propensity	11
1.6 Probability is not What Everyone Believes	12
1.7 Probability is not a Syntactic Relation	15
1.8 Probabilities are not Subjective	21
Chapter 2 What Degrees of Belief Aren't	35
2.1 Analyses of Degree of Belief	35
2.2 Why the Betting Analysis of Degrees of Belief Fails	37
2.3 Dutch Book Arguments Fail	41
2.4 Other Critiques of Dutch Book Arguments	48
2.5 The Betting Analysis as Analogy	51
Chapter 3 What Degrees of Belief Are	55
3.1 The Equivalence Analysis	55
3.2 Outline of Chapter	57
3.3 Precise Models	59
3.4 Updating Precise Models	64
3.5 Introducing Imprecision	70
3.6 The Single Model Approach	81

3.7	<i>Many-Urn Models</i>	93
3.8	<i>Updating</i>	99
3.9	<i>Conglomerability and Lower Envelopes</i>	109
3.10	<i>Real-Valued Degrees of Belief</i>	119
Appendix 3A	<i>Proof of Theorem 3.3.1</i>	121
Appendix 3B	<i>Proof of Theorem 3.6.1</i>	123
Appendix 3C	<i>More on Dutch Book Arguments</i>	125
Chapter 4 What Probability Is		132
4.1	<i>The Probability Relation</i>	132
4.2	<i>Necessitarian Probability</i>	136
4.3	<i>Ambiguity and Relations</i>	138
4.4	<i>Isms</i>	150
4.5	<i>Lewis's New Principle</i>	156
Chapter 5 Supervaluations		163
5.1	<i>Introduction</i>	163
5.2	<i>Supervaluations</i>	164
5.3	<i>Probability Sentences</i>	166
5.4	<i>Scope and the T-schema</i>	169
5.5	<i>Validity</i>	179
5.6	<i>Models and Conceptual Truths</i>	182
5.7	<i>Local and General Supervaluationism</i>	189
5.8	<i>The Reasonableness of Imprecision</i>	191
5.9	<i>The Reasonableness of Precision</i>	196
Chapter 6 Objections		199
6.1	<i>Introduction</i>	199
6.2	<i>There are no such things</i>	201

6.3	<i>Probability Relations Between Simple Propositions</i>	202
6.4	<i>The Sorites Objection</i>	204
6.5	<i>The Probability Calculus</i>	205
6.6	<i>The Principle of Indifference</i>	206
6.7	<i>Uncertain Evidence</i>	208
6.8	<i>A Recent Addition - Dependence on A Priori Assumptions</i>	210
Chapter 7 Philosophical Predecessors		214
7.1	<i>Levi</i>	214
7.2	<i>van Fraassen</i>	221
7.3	<i>Jeffrey</i>	226
7.4	<i>Kyburg</i>	228
Chapter 8 Constructivist Probability		238
8.1	<i>Motivations for a Constructivist Approach to Probability</i>	238
8.2	<i>The Morgan - Leblanc - Mares Calculus</i>	242
8.3	<i>Developing a Constructivist Probability Calculus</i>	244
8.4	<i>Kripke Trees</i>	246
8.5	<i>Intuitionist Probability</i>	252
8.6	<i>Updating</i>	255
8.7	<i>Objections</i>	259
<i>Appendix 8A Proof of Soundness of the Axioms</i>		262
Chapter 9 Vague Decision Theory		265
9.1	<i>Introduction</i>	265
9.2	<i>Unstructured Decision Theories</i>	270
9.3	<i>Levi's Rule</i>	274
9.4	<i>Conservatism</i>	280
9.5	<i>Caprice</i>	283

9.6	<i>Arguments For Caprice</i>	287
9.7	<i>Monte Hall Again</i>	289
Chapter 10 Keynes and Probability		291
10.1	<i>Introduction</i>	291
10.2	<i>Keynes's Pure Theory of Probability</i>	292
10.3	<i>Keynes's Applied Theory of Probability</i>	299
10.4	<i>Keynes and Conventions</i>	304
Chapter 11 The Economic Consequences of Uncertainty		311
11.1	<i>Uncertainty, Investment and Unemployment</i>	311
11.2	<i>Two Consequences of Uncertainty</i>	314
11.3	<i>Uncertainty and Money</i>	317
11.4	<i>Uncertainty and Liquidity Preference</i>	321
11.5	<i>Uncertainty and Indecision</i>	328
11.6	<i>Disquietude</i>	331
11.7	<i>Summary</i>	334
Bibliography		336
Name Index		351
Subject Index		354

Synopsis

This dissertation looks at a set of interconnected questions concerning the foundations of probability, and gives a series of interconnected answers. At its core is a piece of old-fashioned philosophical analysis, working out what probability is. Or equivalently, investigating the semantic question of what is the meaning of ‘probability’? Like Keynes and Carnap, I say that probability is degree of reasonable belief. This immediately raises an epistemological question, which degrees count as reasonable? To solve that in its full generality would mean the end of human inquiry, so that won’t be attempted here. Rather I will follow tradition and merely investigate which sets of partial beliefs are coherent.

The standard answer to this question, what is commonly called the Bayesian answer, says that degrees of belief are coherent iff they form a probability function. I disagree with the way this is usually justified, but subject to an important qualification I accept the answer. The important qualification is that degrees of belief may be imprecise, or vague.

Part one of the dissertation, chapters 1 to 6, looks largely at the consequences of this qualification for the semantic and epistemological questions already mentioned. It turns out that when we allow degrees of belief to be imprecise, we can discharge potentially fatal objections to some philosophically attractive theses. Two of these, that probability is degree of reasonable belief and that the probability calculus provides coherence constraints on partial beliefs, have been mentioned. Others include the claim, defended in chapter 4, that chance is probability given total history.

As well as these semantic and epistemological questions, studies of the foundations of probability usually include a detailed discussion of decision theory. For reasons set out in chapter 2, I deny we can gain epistemological insights from decision theory. Nevertheless, it is an interesting field to study on its own, and it might be expected that there would be decision theoretic consequences

of allowing imprecise degrees of belief. As I show in part two, this expectation seems to be mistaken. Chapter 9 shows that there aren't interesting consequences of this theory for decision theory proper, and chapters 10 and 11 show that Keynes's attempt to use imprecision in degrees of belief to derive a distinctive theory of interest rates is unsound.

Chapters 7 and 8 provide a link between these two parts. In chapter 7 I look at some previous philosophical investigations into the effects of imprecision. In chapter 8 I develop what I take to be the best competitor to the theory defended here – a constructivist theory of probability. On this view degrees of belief are precise, but the relevant coherence constraint is a constructivist probability calculus. This view is, I think, mistaken, but the calculus has some intrinsic interest, and there are at least enough arguments for it to warrant a chapter-length examination.

Declaration of Originality

- f* This dissertation contains no material which has been submitted for the award of any other degree or diploma in any university or other institution.
- f* To the best of my knowledge, this dissertation contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

Signed: Date:

Brian Weatherson

Acknowledgments

As with any doctoral dissertation, this could not have been completed without supervisory support and assistance. My supervisors, Graham Oppy and Richard Holton, have excised many of the bad ideas from the dissertation and assisted greatly in the development of the good ones. The usual disclaimer about remaining mistakes applies. Graham's diligence in detecting spelling and grammatical errors was probably all that prevented an embarrassing surfeit of errors remaining herein.

I also owe a great debt to the departments I attended while undertaking the degree: the Philosophy department at Monash University and the Philosophy Program in RISS, Australian National University. In particular, I would like to thank John Bigelow, Lloyd Humberstone, Philip Pettit, Daniel Nolan, James Chase and Konrad Talmont-Kaminski for discussions on subjects in my dissertation, as well as the audiences at several seminars I presented in those departments over the last four years.

Chapter 1

What Probability Isn't

1.1 Introduction

Part one of this dissertation defends the view that we should analyse probability as reasonable degree of belief. So the correct analysis of the sentence *The probability that Oswald killed JFK is greater than 0.5* is that the only degrees of belief in *Oswald killed JFK* that are reasonable are greater than 0.5. This will obviously have to be relativised to some evidence; the sentence is not refuted by the existence of people who have never heard of Oswald or Kennedy and who can thus reasonably refrain from having a high degree of belief in *Oswald killed JFK*. Hence I claim that probability sentences contain an elliptical reference to evidence; in chapter 4 I'll say more about how this reference works. That probability sentences are in part elliptical is not at all controversial: virtually every theory of probability does this in some way.

This approach to probability puts my account in the tradition of Keynes (1921a) and Carnap (1950). They advocated this analysis of probability, and what I regard as one of its corollaries, that probability sentences are non-contingent. I differ from these theorists in one important respect. They thought that probability is a 'logical' concept, in a rather narrow sense. That is, they thought that the truth of probability sentences could be deduced from their syntactical structure in an ideal language. As I will argue in section 1.7, there are pragmatic and theoretical reasons for rejecting this approach. Nevertheless, there is no reason why all non-contingent truths must be true in virtue of syntactic form, so I can differ from Keynes and Carnap on this point while holding on to their more important insight.

Probability is not just used in sentences like *The probability that Oswald killed Kennedy is α* . We also use probability in a purely mathematical sense when we talk about the probability calculus. Part of my theory of probability is to explain the connection between these two facets of probability. In chapter 2 I argue that a very common way of proving a connection (the 'Dutch

Book argument') is unsound, however in chapter 3 I give a new argument for this connection. The argument only works on the assumption that degrees of belief ought be precise, or what is equivalent, completely ordered, and that assumption is false. The argument is, however, illuminating as to what we ought say about situations where degrees of belief are partially ordered. The probability calculus will be so important to what follows that I give a brief introduction to it in section 2.

The rest of this chapter is to set out the common objections to all the other analyses of probability on the market. The aim is not to provide a conclusive refutation, but as Ramsey said to “show that [they are] not so completely satisfactory as to render futile any attempt to treat the subject from a rather different point of view” (Ramsey 1926: 166). The most important objection to some of these theories is provided by the rest of the dissertation. In particular those theories of probability which are defended by showing necessitarian analyses to be implausible are weakened not so much by my direct attacks on them as by my defence of their rival.

Sections 3 argues that probability should not be analysed as actual frequency, and section 4 shows that analysing probability as modal frequency is no better. In section 5 I argue that Popper's conception of probability as propensity cannot explain probability sentences about past events, and hence cannot be a complete theory. Section 6 argues that we cannot adopt Ayer's conventionalist approach to probability, because of the problem of unknown conventions. Section 7 looks at the problems with the syntactic theories of probability defended by Keynes and Carnap. Finally, and perhaps most importantly, in section 8 I examine the various kinds of theory of probability called 'subjective'. This includes the necessitarian theory defended here. Following Carnap I argue that this theory is not properly called subjectivist, and following Ayer I argue that most other theories called subjectivist are flawed.

A note on notation before I start. I will often talk about probability sentences; indeed the overall project here could be described as trying to analyse probability sentences. There are more types of probability sentence than I have indicated above. I also intend the term to refer to sentences like the following:

- Oswald probably killed JFK.
- It's more probable that Oswald killed JFK than that O. J. Simpson killed his wife.
- The probability that Oswald killed JFK is 0.6.
- The probability that Oswald killed JFK given the forensic evidence is 0.6.

That is, I take probability sentences to come in qualitative, comparative, quantitative forms as well as in conditional and unconditional forms. Of course I could by mixing these forms come up with even more examples, but I hope this is enough to indicate the field in which I'm interested.

1.2 The Probability Calculus

Mathematically, probability functions have as their domain a field of sets, and as their range reals in $[0, 1]$. However, the probability sentences that I'm taking to be the explicandum of our theory seem to refer to the probability of an event or sentence. We solve this little problem by saying that probability sentences talk about the probability of propositions, and propositions are just sets of possible worlds. The proposition p is just the set of possible worlds in which p . Hence we can interpret the 'universe' in the mathematical representation as the set of all possible worlds, and the field as a set of propositions.

We take as given a possibility space U , and a field F of subsets of U . That F is a field just means it includes U , and is closed under complementation, union and intersection. $Pr: F \rightarrow [0, 1]$ is a simple probability function just in case it satisfies the following three axioms.

(Pr1) For all $A \in F$, $Pr(A) \geq 0$;

(Pr2) $Pr(U) = 1$

(Pr3) If $A, B \in F$ and $A \cap B = \emptyset$ then $Pr(A) + Pr(B) = Pr(A \cup B)$

In propositional terms, (Pr2) says that the probability of any (classical) tautology is 1, and (Pr3) says that if p and q are inconsistent then the probability of $p \vee q$ is the probability of p plus the probability of q . The canonical statement of all this is in Kolmogorov (1933). He makes two complications to the theory. The first is to extend it to conditional probability functions. Often the axiomatisations for conditional probability functions are given in such a way that probability could be conditional or non-conditional. I think it's neater to only allow conditional probabilities, and since I think all probability sentences make elliptical (or explicit) reference to evidence, there is a philosophical justification for this. So the axiomatisation for conditional probability functions $Pr: F \times F \rightarrow U$ is as follows.

For all $A, B, C \in F$

(CP1) $Pr(A | B) \geq 0$

(CP2) $Pr(U | A) = 1$

(CP3) If $A, B, C \in F$ and $A \cap B = \emptyset$ then $Pr(A | C) + Pr(B | C) = Pr((A \cup B) | C)$

(CP4) $Pr(A | B \& C) \cdot Pr(B | C) = Pr(A \& B | C)$

The notation $Pr(A | B)$ is read as ‘the probability of A given B ’. We can recover the ‘unconditional’ probability $Pr(A)$ as $Pr(A | U)$. When the simplification is harmless and aids the exposition I will occasionally talk about simple, or unconditional, probability functions, but the main focus will be on analysing probability sentences by using of conditional probability functions. Note that we can almost recover a conditional probability function from a simple one by setting $Pr(A | B) =_{df} Pr(A \& B) / Pr(B)$. The problem is that this definition fails when $Pr(B) = 0$. It seems on the whole simpler to take conditional probability functions as basic.¹

¹ It has been reported to me that Alan Hájek's as yet unpublished Ph.D. thesis contains a wide range of

The other complication Kolmogorov makes is to extend the additivity axiom, (Pr3) or (CP3), from a principle of ‘finite additivity’ to one of ‘countable additivity’. This involves the adoption of a new axiom, (CP5).

(CP5) If A_1, \dots, A_n, \dots are pairwise disjoint elements of F , then $\bullet \Pr(A_i|C) \bullet \Pr(\bullet A_i|C)$

It is hardly ever suggested that this be extended to cases where there are more than denumerably many A 's, for example where there is one element of the A 's for every real in $[0, 1]$.² However, there is some debate about whether even extension to the countable case is plausible. Kolmogorov merely defended it on grounds of mathematical convenience, which is hardly telling. The following example shows both how (CP5) is independent of the other axioms, and why we might not want this axiom.

Say we know x is a natural number, but have no idea about which natural number it is. In this case we might think it appropriate to spread the probability evenly over every element of N . That is, for any natural number n , set $\Pr(x = n) = 0$. Or in conditional language, set $\Pr(x = n | x \in N) = 0$. Now this is clearly consistent with the axioms apart from (CP5), and it is clearly inconsistent with (CP5). To see this, set A_i as $x = i$ for all i . The probability of each A_i is 0, but the probability of their union, $x \in N$, is 1. de Finetti thought this probability function was so obviously reasonable in the circumstances that he rejected Kolmogorov's axiom (de Finetti 1974: 121). In chapter 3 I'll look into this in more detail, but to get ahead of myself a little bit, I don't think (CP5) has been proven to be appropriate for our usage of probability. And since I think there's a burden of proof

arguments for this conclusion, including arguments against resolving the difficulty of undefined conditional probabilities by moving to infinitesimals. However, without having seen that thesis, I am unable to comment in any detail on it.

² Though Lewis (1994) seems to suggest that we might need such ‘strong forms of additivity’ to deal with the infinitesimal-valued probabilities he posits.

on the proponent of a new axiom, for now I take de Finetti's side of this debate. However, I'm not as convinced as de Finetti that there will never be an argument for countable additivity.

It might be worth noting one of the confusing nomenclatures in this field, if just to note that I won't be adopting it. Sometimes the term 'finitely additive' is used for only those probability functions which do not satisfy countable additivity, our (CP5). This is misleading because of course countably additive functions are also finitely additive on the most natural interpretation of that term. That is, they satisfy (CP1) to (CP4). When I use the term 'finitely additive' that is precisely what I will mean, but to minimise confusion I'll just try not to use it at all.

So all I mean by a probability function is something satisfying (CP1) to (CP4). These will be important to our eventual analysis of probability sentences, but for now we can leave mathematics and return to philosophical analysis. Or at least to refuting philosophical analyses.

1.3 Probability is not Frequency

It could be the case that *The probability that Oswald killed JFK is more than 0.5*, which we'll abbreviate to O , is true even if Oswald did not kill JFK. Since there was only one JFK assassination, that would make Oswald's frequency of being JFK's assassin 0. Yet this wouldn't, one suspects, make us say that O is necessarily false now. If there is a lot of evidence for Oswald's guilt, as many people seem to believe, then O will be true. Hence we cannot interpret O as a statement about frequencies, in this simple sense.

Nor could probability be long-run frequency. If probability is long-run frequency it must be that Oswald being the assassin is one type of event, and the JFK assassination is another, and the ratio of events of the first type amongst events of the second is more than 0.5. Now on the one hand if we specify these types too closely then we will be back to the problem that there is at most one event of each type, so the ratio will be 0 or 1. On the other hand if we specify too coarsely, we lose any theoretical motivation for linking probability and frequency.

Consider, for example, some of the possible event types E_1 and E_2 such that Oswald being the assassin is an instance of E_1 and the assassination is an instance of E_2 and the probability of Oswald being the assassin is the frequency of E_1 events amongst the E_2 . (I.e. $n(E_1 \& E_2) / n(E_2)$ or some limit of this, where $n(E)$ is the number of times E occurs). If E_1 is Oswald being the assassin and E_2 is there being an assassination, then O will be obviously false, but presumably it could be true. If E_1 is the initial suspect being guilty then the probability of initial suspects being guilty at every assassination will be constant, which seems mistaken. Similar considerations preclude E_1 being say, a communist sympathiser is guilty, or being that someone who killed someone else on the day of the assassination is the killer. If we start taking conjunctions of these, say E_1 being the initial suspect, who is a communist sympathiser and a known killer, is guilty and E_2 is that there is an assassination where the initial suspect is a communist sympathiser and a known killer we risk the classes contracting to size 1 again, and the frequencies hence being either 0 or 1.

Even when the frequency analysis gives the correct output, it seems to get the direction of explanation wrong. Moving from assassinations to casinos, let E_2 be the event that a standard (i.e. 37-slot) roulette wheel is spun and E_1 the event that the ball lands in 1. The probability of the ball landing 1 is $1/37$, which is presumably also the frequency. I have just made a well-balanced, apparently fair 35-slot roulette wheel. The probability of the ball landing 1 on first spin is, it would seem, $1/35$, even if this is the first ever spin of a 35-slot roulette wheel, and indeed even if it is the only ever spin of such a wheel. This is just the point of the previous argument, however there is a larger problem for the frequency analysis.

Say that a schmoulette wheel is a 35-slot roulette wheel made today or a 37-slot wheel made any other day. Let E_2 be the event that a schmoulette wheel is spun, and E_1 the event that the ball lands 1. The frequency of E_1 amongst the E_2 will in the long run be little different from $1/37$. This doesn't alter the fact that the probability that the ball will land 1 on the first spin of my 35-slot


wheel is $1/35$, even though the spin is an event of type E_2 and the ball landing 1 of type E_1 . The conclusion I draw is that the events in E_1 and E_2 must be homogenous in some way if the frequency analysis is to give the correct response. However, I suspect there will be no way of defining this homogeneity except by reference to probability. In other words, it seems that it must be probability that determines frequency, rather than frequency determining probability. Unless we already know the probability of particular events we can't determine appropriate event types, and without that we can't determine the frequency of a type of event³.

Russell (1948: 384) showed, when the cardinality of E_1 and E_2 is infinite, the ratio of occurrences of E_1 to $\neg E_1$ amongst the E_2 can only be defined as a limit. That is, we list the occurrences of E_2 in some order, and say the frequency of E_1 is the limit as n tends to infinity of the number of events in the first n which are E_1 to n . However, the limit of this ratio depends not only on the membership of E_2 , but on its ordering. For example, if we order the natural numbers in the standard way (i.e. 1, 2, 3, ...) then as n tends to infinity the ratio of the number of primes less than or equal to n to n will tend to 0. So the long-run frequency of primes in the natural numbers is 0. However, if we simply re-order the numbers, we can make this limit be $1/2$, or 1, or indeed any number we care to choose in $[0, 1]$. So frequency can't be defined as a relation between classes, but only as a ratio between sequences. As Russell remarked, "This seems strange." (1948: 385).

There are also a multitude of theoretical reasons for not equating frequency and probability. One is what van Fraassen (1989) calls the *horizontal-vertical* problem. Assuming we're trying to work out the probability of an event that will (or will not) happen in 2000, say a Democrat winning the 2000 U. S. Presidential election. Consider a branching-time model of the universe, with the

³ In section 7.4 I argue against Kyburg's attempt to define homogeneity in just this way. Kyburg is not a frequency theorist; he is a logical theorist who thinks that probability refers to a metalinguistic relation between a sentence and its evidence whose value is determined by the most pertinent statement about frequencies in the evidence.

possible world time-slices being points on a Cartesian plane, and with actual time as the y -axis. This is drawn in the diagram below. Possible worlds are functions $x = f(y)$. In the diagram below the actual world @ (the bold line) is represented by the function $x = c$. The dotted horizontal lines then represent all the ways the world could be at various points in time. The large dot is the way the actual world is now. The other lines leaving this world represent worlds which have the same past as ours, but diverge between now and the year 2000. There are of course infinitely many such worlds, but only finitely many can be drawn.



To work out the relative frequency of one event type given another, we only have to look at @. In particular, we look at the vertical line $x = c$, and work out the ratio of points on it that are of type E_1 & E_2 to those that are of type E_2 . If this is impossible we work out the limit of this ratio as time tends to infinity. However, to work out the probability now of a certain event happening in 2000, we presumably have to look at the ratio of points on $t = 2000$ which are of that type to those that are not, or more likely some weighted average of this type, or more likely again a limit of some such weighted average. The important point is that what is important to the probability of a Democrat winning is a ratio of some kind on the horizontal line $t = 2000$, not on the vertical line @. Frequencies measure the wrong things to be probabilities.

Finally, there is the problem Kyburg (1961: 22) noted about the applicability of the frequency analysis. Let's take a case where the frequentist should be on solid ground, the case where we try to work out the probability of a coin toss landing heads. Coin tosses happen often enough, and are homogenous enough, that at least some of the standard objections to frequentism are irrelevant. Perhaps then the frequentist can explain what we mean by 'The probability of a coin toss landing heads is $1/2$ '. However, as Kyburg points out by their own lights we cannot mean anything by 'The probability of the next coin toss landing heads is $1/2$ '. The frequentist can only talk about the probability of events which are outcomes of trials repeated very often, perhaps infinitely. However, 'the next coin toss' is not a repeated trial, hence they can't talk about the probability of it. So even in cases where they appear most comfortable, the frequentist only gets away with their story by shifting from a definite to an indefinite article.

1.4 Probability is not Modal Frequency

Recognising the horizontal-vertical problem, some people have argued that probability is modal frequency. That is, the probability of p is the frequency of p across the possible worlds, or the ratio of p -worlds to all worlds. This does solve the horizontal-vertical problem, and it solves the problem of one-off events (like the JFK assassination) having a probability. But it seems to make a fundamental mistake about the nature of the possible worlds. There are just too many of them for this to get off the ground. Provided p is contingent, there are infinitely many worlds in which p , and infinitely many in which $\neg p$. Now there are ways to get around this, indeed my theory could be considered such a way, but when we take it we seem to not have a modal frequency theory. Moreover, it is hard to see how on this analysis we could think the probability of a Democrat winning in 2000 is higher than that of a Republican winning. It's not that there are more worlds in which Democrats go on to win than in which Republicans do, just that (at present) the Democrat-winning worlds are more probable.

We might hold that probability is modal frequency among the accessible worlds, or that it is some kind of weighted modal frequency. I have no objection to such a view; indeed it is quite similar to a view I adopt. However, adopting this as an analysis seems to me to get the order of explanation wrong. The frequency of a highly probable event among accessible worlds is high simply because the event is probable. That is, the accessible worlds are accessible because they are probable, they are not probable because they are accessible. So I think such analyses may be extensionally correct, but even if they are will be flawed as analyses because their ‘direction of fit’ is wrong.

1.5 Probability is not Propensity

Popper (1959) held that probability should be analysed as propensity. This seems to make sense when we are looking to analyse probability statements in, say, quantum mechanics. But it doesn’t make sense in a number of senses in which probability is used. In particular, it doesn’t seem to work when we are considering the probability of events which, if they did happen, would have happened in the past, or the probability of laws of nature.

As an example of the first type of problem, note that it makes sense to talk about the probability that Oswald is guilty. Now this doesn’t mean (and nor would Popper have said it meant) that Oswald is likely to commit more crimes. Even if we are now totally convinced that Oswald’s current propensity to commit crimes is low (because he’s dead), the probability that he was a killer can be high. So we can at most talk about what the propensity was. Even this seems implausible, as the following is not contradictory: “It seems highly probable on the basis of the forensic evidence that Oswald did it, though it would have been completely out of character for him”. Assuming Oswald’s character determines his propensity to commit crimes, this means we can distinguish between probability and past propensity. So it must be current propensity that matters. But the current propensity must be either 0 or 1. The world is already either an ‘Oswald-did-it’ world or an ‘Oswald-didn’t-do-it’ world, so its propensity to become one of these

is 0 or 1. Yet the probability that Oswald did it on the basis of a certain body of evidence can be between 0 and 1.

As an example of the second type of problem, note that it seems plausible, at least when doing historical reconstructions, to talk about the probability that the laws are one way rather than another. We can talk sensibly about a certain experiment making one theory more or less probable. But we can't, it would seem, make any sense of propensity statements without assuming laws as given. What is usually referred to as the propensity of, say, atoms to decay is at best a matter of natural law. If we take the law as up for question, the propensity is indeterminate. Since Popper does not think that all probabilities are indeterminate, it must be that we take laws as given when determining probabilities. Hence the probability of any actual law must be 1, and the probability of any counterlegal is 0. But this goes against our evidence that the probability of a purported law can change with experiments. So the propensity analysis must fail. These two counterexamples can be connected. The propensity theory cannot explain statements like 'On the evidence the ancient Egyptians had, it was highly probable that the earth was flat, but on the evidence we have today, this is highly improbable' both because it discusses the probability of prior events and it allows for counterlegals to have a positive probability.

1.6 Probability is not What Everyone Believes

Before getting onto orthodox subjectivist analyses in section 8, I want to address here conventionalist theories of probability, which are quite similar. The conventionalist, or intersubjectivist, argues that the probability of p given some evidence q is the degree of belief which the community holds to be appropriate in p given that evidence. On some tellings, the conventionalist agrees with the necessitarian position advocated here that probability should be analysed in terms of reasonable degrees of belief. They even proffer a broadly realist conception of what is reasonable. However, that conception is so different to what I am defending that it amounts to a different analysis.

The conventionalist account is historically important because it seems to be the theory of probability in Ayer's *Language, Truth and Logic*. I say 'seems' because Ayer isn't particularly explicit on this point, and the discussion amounts to no more than a couple of pages. The main evidence is the following quotes.

To say that an observation increases the probability of a hypothesis ... is equivalent to saying that the observation increases the degree of confidence with which it is rational to entertain the hypothesis. And here we may repeat that the rationality of a belief is defined, not by reference to any absolute standard, but by reference to our own actual practice.

[W]hen a man relates belief to observation in a way which is inconsistent with the accredited scientific method of evaluating hypotheses ... he is mistaken about the probability of the propositions which he believes (Ayer 1936: 100-101).

More recently various writers such as Gillies (1988, 1991), Runde (1994a), Davis (1994) and Bateman (1996) have held that Keynes moved to a conventionalist position when he wrote his later economics. Some of these writers, particularly Gillies and Runde, seem to endorse this shift.

It is a little surprising at first that Ayer takes this position on probability. I expected Ayer to adopt a position that was both subjectivist and non-cognitivist about probability, much as he does about ethics. The reason he does not do this is two-fold. First, he recognised some of the good objections to subjectivism, at least as it is commonly presented. Secondly, his verification principle is expressed in terms of probability. A sentence is meaningful, says Ayer, iff it is verifiable. But to make this plausible we have to adopt what Ayer calls the 'weak' conception of verifiability. "[A proposition] is verifiable, in the weak sense, if it is possible for experience to render it probable"

(Ayer 1936: 37). Now if what was probable varied from person to person (as some subjectivists assert) it would turn out that which sentences were meaningful varied from person to person. This is much too implausible for Ayer. Alternatively, if we go fully expressivist (or non-cognitivist) about probability, and say that there is no fact of the matter as to whether or not a proposition has been rendered probable, there will be no fact of the matter as to whether some sentences are verifiable. This is again not a conclusion Ayer wants.

However, the conventionalist move has problems of its own. There is one argument against it which seems quite powerful to me, but which is obviously question-begging. On Ayer's story whether q renders p probable will depend not just upon p and q , and perhaps on background facts, but on the prevailing scientific standards. Probability sentences for Ayer presumably have an elliptical reference to these standards. Now this seems completely implausible, but since Ayer happily accepts it we can hardly urge it as an argument. I only mention it to remind the reader that their intuitions on this matter may differ from Ayer's.

The more substantial problem for Ayer is what I call the problem of unknown conventions. Since it is an empirical fact that convention A is operative in our society, rather than say convention B , this is only something we can learn by experience. That is, we learn it because we acquire evidence for it. Presumably this evidence, like all other evidence, could be misleading. So say an agent has evidence q , and that evidence provides strong but misleading support for the proposition that convention B is operative. Now assume, as again seems possible, that conventions A and B provide different directions as to the appropriate degree of belief in p given q . Say A says p ought be believed to degree 0.3, and B says it should be believed to degree 0.8. Now, what ought our agent do?

The conventionalist says that the agent ought believe p to degree 0.3. Note, however, that if the agent does the best they can do to accord with the conventions, that is, arrange their beliefs in

accord with what they reasonably believe to be the conventions, they will believe p to degree 0.8. I don't see how the conventionalist can criticise an agent who does follow convention B in these circumstances. After all, that agent has done what they could to satisfy conventionalist doctrines; they have arranged their beliefs in accord with what they have reasonably taken to be the conventions of society. So I think the conventionalist is forced to say this kind of person both is and is not reasonable.

We can make the same point in a more dramatic way. The conventions in which we are interested are just rules for converting evidence to reasonable degrees of belief. If someone believes all the conversions, they believe the convention, even if they can't express it. And plausibly we can analyse believing a particular conversion as being disposed to make it in the right circumstances. That is, if an agent is disposed to believe p to degree x on evidence q , they believe that the relevant rule converts evidence q to degree of belief x in p . They might, in some circumstances, not know that they believe it, but believe it they do. The conventionalist says reasonable agents will always have these dispositions. Hence all reasonable agents will believe the conventions are what they actually are, even on no evidence whatsoever. Since, as was noted, what the conventions are is for Ayer an empirical fact, he imposes upon his rational agents a requirement to believe an empirical fact on no evidence at all. This is hardly plausible, so Ayer's conventionalism about probability fails.

1.7 Probability is not a Syntactic Relation

In his (1921a) Keynes argued that probability referred to 'partial entailment' relationships, of which classical entailment is merely a limiting case. The spirit of this approach was adopted by Carnap in his (1950) and subsequent works. There are, as I have noted, strong similarities between my approach and the Keynes-Carnap approach. However, there are two crucial points on which I differ from Carnap, and the point of this section is to briefly set out Carnap's theory and my grounds for dissenting from it.

The exposition of Carnap's position given here largely follows the exposition in T. Fine (1973: Ch. 7) and Carnap's own summary in his (1963). In simple terms, Carnap analyses probability in terms of degree of belief. The probability of h given e is the degree of rational belief in h given e . That's entirely accurate, but Carnap didn't like it as a description because it might have misleadingly subjectivist connotations. So as a next approximation he said the probability of h given e is the degree of confirmation of h by e . But this term too could be misinterpreted, as his exchange with Popper in the 1950s indicated. So he eventually defined the probability of h given e as the 'rational subjective value' in utils of a bet which pays 1 util if h and nothing otherwise to an agent with evidence e .

It isn't clear why Carnap thinks this explanation of probability should imply it is always numerically valued. The comments at (1963: 972) suggest he thinks this is necessary for probability to be used in rational decision making. In any case, he set himself the task of developing a quantitative theory of probability in account with the above analysis. Carnap thinks, correctly in my opinion, that the concept of probability⁴ he is working with is crucial to induction. Sound inductions are those where the logical probability of the conclusion given the premises is high. So he draws the following conclusions:

- (a) The reasons [for accepting axioms of inductive logic] are based upon our intuitive judgements concerning inductive validity, i.e. concerning inductive rationality of practical decisions (e.g. about bets); therefore:

⁴ Actually Carnap thought there were two concepts of probability, one based on 'logical probability' and the other based on frequency, so the text might be a bit misleading here. However, it was the logical concept, his probability₁ which attracted most attention, and which in later writings he referred to simply as probability. Hence most commentators have adopted the convention I'm using of referring to probability₁ as Carnap's conception of probability.

- (b) It is impossible to give a purely deductive justification of induction.
- (c) The reasons are *a priori* (Carnap 1963: 978).

I think (c) is correct, but Carnap goes further. He has taken h and e to be sentences, not propositions, and he thinks that the probability of h given e , like the provability of h from e , can be determined by purely syntactic considerations. The position I will take is that while probability sentences are non-contingent, they are like ‘All bachelors are unmarried’ in being true in virtue of their non-syntactic features.

To spell out this qualitative concept of logical probability, Carnap attempts to develop a c -function which will give the value for the ‘degree of confirmation’ of h given e for any h, e in a given language, written as $c(h, e)$. To narrow down the class of functions which could serve the role of c , he adopts a number of axioms. These fall into three categories. The first category does enough to say c is a conditional probability function. The second are symmetry constraints. So for example we have an axiom saying that universal substitution of one name for another throughout h and e leaves the value of $c(h, e)$ unchanged. And similarly universally substituting one predicate for another from the same family⁵, or substituting one family of predicates for another family with the same number of elements leaves $c(h, e)$ unchanged. Finally, adding new families of predicates to the language will leave $c(h, e)$ unchanged, as will adding terms for new individuals, provided h and e contain no quantifiers. It is these invariance postulates which prompt me to describe Carnap’s as a ‘syntactic’ theory of probability. Finally, Carnap has three axioms asserting that c must allow an agent to learn from experience. However, as Fine shows these don’t do much to restrict the class of permissible c -functions, and in some cases are simply redundant.

⁵ A family of predicates is a set of predicates such that every possible individual satisfies exactly one of them.

I will note two types of objection to Carnap's approach. The first essentially object to his claim to have developed a quantitative theory; the alleged faults are caused by his claim that c is real-valued. The second object to the claim that probability is syntactic. I think both classes of objection will fail.

There are four problems with Carnap's claim to have developed a quantitative account. van Fraassen (1989: 119-25) stresses the point that despite Carnap's aim, the axioms he gives do not suffice to specify a unique c -function. Carnap of course knew that we had to posit a continuum of c -functions, but could say little about how to choose between them (Carnap 1952). Fine notes that if we make a seemingly plausible extension of Carnap's axioms, if we insist that uniform substitutions of one complete description of the world for another leaves probability unchanged, we are led into inconsistency. From this we conclude that not all symmetry requirements are met, that Carnap's axioms aren't as plausible as seemed at first (since the intuitions which grounded them perhaps provide equal support to inconsistent axioms) and, Fine argues, that Carnap must say that the probability of h given e is not determined by the meaning of h and e . The point here is that if the language includes the family of predicates $\{red, not-red\}$ then the probability of a is $not-red$ given a tautology is $1/2$, whereas if it includes the family $\{dark red, light red, not-red\}$ the probability of a is $not-red$ given a tautology is $1/3$. Thirdly, as Howson and Urbach (1989: Ch. 3) urge, not all symmetry requirements can be met at once. They note that different symmetry requirements are inconsistent. Fourthly, as Keynes (1921a) notes, there is no principled way to avoid the paradoxes of indifference if we insist that all probabilities are numerically valued.

Some of these problems look like they'll go away if we allow there to be more than one permissible c -function. Carnap at one point (1963: 971) goes very close to endorsing just this. However, there are a separate set of objections which can be levelled at the syntactic parts of Carnap's theory. There are two objections which can be levelled at this, the first based around the

problem of non-projectability and the second around some objections of Jeffrey to Carnap's theory of evidence.

For our purposes it will be preferable to use the discussion of non-projectability in Russell (1948) rather than the more standard discussion in Goodman.⁶ As Russell notes, for many predicates F and G , the inference 'All F s observed so far have been G s' therefore 'The probability that all F s are G s is high' is sound. Or again, the probability that all F s are G s given all F s observed so far have been G s must be high. If we are to base probability around syntactic considerations and get started at all, we will have to accept this rule. However, as Russell also notes, some inferences of this form are clearly unsound. If a farmer has only seen cows in Heresfordshire so far in his life, this is no justification for believing that probably all cows are in Heresfordshire. That is, the inference is unsound when F is 'is a cow' and G is 'is in Heresfordshire'. But the unsound inference has the same syntactic form as some sound inferences. So probability can't be based on syntactic form.

Perhaps there is a way out of this problem. We could, for example, restrict the language in which we allow inferences to be made, so there is no way in the language to represent the troublesome

⁶ Discovery of the problems for induction caused by non-projectability is usually credited to Goodman (1954), or perhaps his (1947). However, the discussion in the earlier paper is rather brief, indeed just confined to the predicate P of an artificial language. Further the problem of non-projectability is urged more as a problem for the analysis of counterfactuals rather than for induction, as is now standard. So for these reasons I prefer giving credit to Russell. It would be interesting to discover when Russell discovered this problem. In his (1940) he is obviously ignorant of it. In the introduction to (1948) he mentions that parts of it are based on lectures he gave in 1944-45, but doesn't make clear which parts. And the papers published from this time in his *Collected Papers* yield little light on the matter. My crediting Russell with this discovery is not meant to say Goodman's book was anything less than an independent discovery, and of course it moved the debate forward and promoted the idea of non-projectability in a way which in the long run proved more effective.

inference as being of the same syntactic form as the sound inferences. This is what Kyburg does in his logical approach. Recently Tooley (1987) has argued that if we are realist about universals, we can restrict the predicates of our canonical language to those universals which exist, and presumably all the universals are projectible. If the existence of universals is non-contingent this might do the work required, though if not the probability sentences will not be *a priori* as Carnap required.

There is, however, a bigger problem. Despite what may be inferred from some philosophy texts, we don't just make inductive inferences or use probability sentences in physical sciences. The idea behind Kyburg's and Tooley's approach is that the ideal language of science will not include troublesome predicates like 'is in Heresfordshire' or Goodman's 'grue'. While it *might* be true that no absolute positional predicates are needed in physical science⁷, this just isn't true in social sciences. At the very least we are going to need predicates like 'in a city', 'in the country' to do the most primitive sociology. So the ideal language of science generally will most likely include predicates which are not projectible.⁸ This isn't yet an argument for saying that we will always end up with gruesome predicates, just an argument for saying that quite a lot more needs to be done to show we are rid of them.

In any case, there's another problem for the syntactic account. For this account to be plausible, we have to be able to specify the evidence we have for a proposition in a finite sentence of some language. But this seems implausible, as Ramsey (1926a) and Jeffrey (1991) have stressed, because of vague evidence. If we view probability as a semantic relationship between

⁷ Clearly predicates referring to relative positions are needed.

⁸ Could we have different languages for social sciences and physical sciences? It seems like a pretty desperate move. It would be hard then to explain sentences which referred to terms from both physical and social sciences, like 'It's more probable that a recession will occur than that this atom will decay in the next n days'.

propositions, rather than a syntactic relationship between sentences, this is no longer a problem, as I outline in chapter 2. The combined effect of these objections to Carnap's account is enough to suggest a different approach could be worthwhile.

1.8 Probabilities are not Subjective

Theories of probability which are called 'subjective', either by their proponents or detractors, abound in the modern literature. Surprisingly then, it is hard to get a clear picture of what is meant by a subjective theory. So my first task in this section is to draw a brief taxonomy of subjective positions. I divide subjective positions into four types, depending on how they deal with two questions. The theories can either define probability in terms of rational degrees of belief or something less, perhaps actual or coherent degrees of belief. And they can be assertoric or expressivist theories. An assertoric theory says that probability sentences make a truth-apt claim about degrees of belief; an expressivist theory says that probability sentences make no claim about how the world is, they just express an attitude. This distinction can be quite clearly seen in looking at different subjectivist ethical theories. An assertoric subjectivist analyses *Torture is wrong* as *Someone (perhaps me) disapproves of torture*. On an expressivist analysis it comes out as *Boo torture!* or some more sophisticated variant on that. There is a fact as to whether I disapprove of torture, hence the assertoric theory is truth-apt, but *Boo torture!* says nothing even plausibly truth-apt. It is, I think, surprising that more subjectivists in probability have not been drawn to expressivist analyses.

One important clarification needs to be made to the above account. Despite some of the quotes I will adduce below, no one seriously believes that probability can be defined purely in terms of actual degrees of belief. If this were the case, there would be no laws of probability at all; as van Fraassen (1990) put it, any such law could be refuted by the existence of a moron. So our moron's degrees of belief have to be made coherent before they can enter into the analysis of probability sentences. As Max Black (1967) put it, degrees of belief have to be at least 'rectified' before we

can use them in analysis. Rectified degrees of belief satisfy some minimal coherence requirements, but nothing more. That these coherence requirements should amount to conformity with the probability calculus is argued for by Dutch Book arguments (see chapter 2) or some variant on them. In non-probabilistic (or classical) epistemology, consistency is not normally considered a sufficient ground for reasonableness. One can consistently believe *The moon is made of green cheese*. Similarly rectified degrees of belief can contain a high degree of belief in *The moon is made of green cheese*. Such a belief state would not be reasonable on any ordinary usage of that term. Despite this, some subjectivists (especially Savage and de Finetti) use reasonable to just mean rectified, and this leads to some confusion. I find Black's terminology clearer, and I'll employ it in what follows.

So we have our four types of subjectivist theory, outlined in the table below.

	Assertoric	Expressivist
Rectified	Type 1 de Finetti, Howson and Urbach	Type 3
Rational	Type 2 Keynes, Carnap	Type 4 Blackburn

The names under each type list adherents of each position. My claim will be that there are strong objections to types 1, 3 and 4 and that type 2 is not properly regarded as subjectivist. The objections to types 1 and 3 are quite old, I will say little about them that is not said by Ayer. My argument that type 2 is not a breed of subjectivism is found entirely in Carnap. And whether or not type 4 subjectivism works seems to turn on whether or not a broadly expressivist program could work, a topic that could cover several chapters on its own. I'll simply note some of the

arguments in the literature as to why it fails. So the originality of this section is confined to its organisation.

Before starting on these objections, I should note one way in which the organisation itself is derivative. Kyburg (1978: 79-80) gives a different four-fold taxonomy of subjectivist positions. The rows are the same as in my table, but the columns refer to a different property. He divides subjectivist theories into theories of decision and theories of degrees of belief. Since I don't regard theories of decision as theories of probability I could hardly adopt this division. Kyburg in turn doesn't consider the distinction between assertoric and expressivist theories. But the motivation for the four-fold taxonomy is in part his paper.

1.8.1 Type 1

It might be thought that type 1 subjectivism is a mere straw man, something I would set up to be knocked down to show the weaknesses of subjectivism. However, it is very hard to read the following quotes as endorsing any other type of subjective theory.

Let us suppose that an individual is obliged to evaluate the rate at which he would be ready to exchange the possession of an arbitrary sum S (positive or negative) dependent on the occurrence of a given event E , for the possession of the sum pS ; we will say by definition that this number p is the measure of the degree of probability attributed by the individual considered to the event E , or, more simply, that p is the probability of E (according to the individual concerned; this specification can be implicit if there's no ambiguity) (de Finetti 1937: 102).

In the personalistic [i.e. subjectivist] concept, probability is an index – in an operational sense to be explained later – of a person’s opinion about an event (Savage 1964: 176).

We shall argue that ... [probabilities] should be understood as subjective assessments of credibility, regulated by the requirement that they be overall consistent. (Howson and Urbach 1989: 39)

Perhaps these might be interpreted as saying that the ordinary language concept of probability is so useless we ought replace it with the concept degree of belief. This might be one interpretation of de Finetti’s later view that “Probability does not exist”, printed in capitals on his (1974: i). However, these quotes seem to be claiming we can analyse *probability* simply as *degree of belief*. And this must be a mistake, because of two arguments from Ayer.

Ayer (1936: 101) rejects this kind of subjectivism about probability because of the ‘obvious objection’ that it doesn’t allow a person to be mistaken about the probability of a proposition. Since the probability of p is just your degree of belief that p , whatever you believe is the probability of p will be its probability. Strictly this mightn’t be quite correct. Presumably a person might believe p to degree 0.2 and believe they believe it to degree 0.3, and hence falsely believe the probability of p is 0.3. So Ayer is wrong to say the subjectivist doesn’t allow mistakes, but they don’t allow mistakes from perfectly introspective agents, which is still implausible.

With this objection Ayer is content to dismiss type 1 subjectivism about probability. There is another objection which we can extract from his dismissal of a simple subjectivist position in metaethics. He dismisses analyses of “ X is wrong” as “I disapprove of X ” by noting that a person can consistently say that they disapprove of things which are not wrong. The equivalent point is a little harder to put in epistemology because of Moore’s paradox, but it can easily be brought out in a little dialogue. If the subjectivist were right, B ’s utterance would be consistent.

A: It is highly probable that the moon is made of green cheese.

B: What *A* says is true, but it is not probable that the moon is made of green cheese.

According to type 1 subjectivism, *A* is making a report about his mental state. *B* can presumably assent to that report, he agrees *A* thinks it probable that the moon is made of green cheese, while consistently saying that *A* is mistaken. But our intuition surely is that *B*'s utterance is inconsistent, which makes type 1 subjectivism implausible.

As an aside, it is possible subjectivists were trying to capture a concept other than *probability*. For instance, in later papers, whenever Savage went to say what the subjectivist (he preferred 'personalist') is claiming, he would give his preferred definition of the *probability for a person* of a proposition. (See for example Savage 1967a and 1967b.) Plausibly he is right *vis a vis* this question, but that concept is not central to probability sentences generally. I am no more making a report about my mental state when I utter *The moon is probably made of green cheese* than when I utter *The moon is made of green cheese*.

1.8.2 Type 2

Type 2 is called subjectivist by Kyburg (1978) who opposes it, and Lewis (1980) who endorses it. Keynes (1921a: 4) vacillates, saying the concept is in part subjectivist, because probability is relative to evidence, and partially not, because it is independent of what anyone thinks. Carnap (1950: 37-50) argues at length that this approach is not properly called subjectivist. I will just repeat some of Carnap's arguments.

Carnap has two primary arguments for calling his probability₁ concept (what I call probability) 'objectivist'. These are that probability sentences are non-contingent and that whether or not they are true is not dependent on anyone's thinking about them. We do use psychologistic terms when giving an analysis of probability, we talk about beliefs, but Carnap has two further reasons for

thinking this doesn't imply subjectivism. First, we never define probability in terms of beliefs *simpliciter*, always in terms of reasonable beliefs. So ours is, in Carnap's language, a *qualified psychologism*. The second reason, which is in part a consequence of the first, is that we can eliminate the psychologistic references from formal presentations. So Carnap defines probability not in terms of degree of reasonable belief, but in what amounts to the same thing, degree of confirmation. Maybe this isn't an improvement, perhaps we can only explain confirmation by reference to reasonable beliefs, but the first two arguments seem sound enough. Hence I think it is possible to define probability in terms of reasonable degrees of belief and oppose subjectivism⁹.

1.8.3 Type 3

The difficulty with looking at possible objections to expressivist interpretations of probability is that there has been so little said about them. This is surprising given the well-known difficulties attending type 1 subjectivism. However, at least this type of expressivist theory seems to do no better. Indeed, it isn't obvious how we avoid the two problems Ayer raises for type 1 subjectivism by saying probability sentences are expressive rather than assertive.

In fact we acquire a new problem. To say that the primary function of a sentence is expressive is no theory at all; we have to say what is being expressed. But it is hard to see how type 3 subjectivism can solve this problem. If we say that what is being expressed is a belief, it looks like probability sentences really are assertoric. After all, the type of utterances that express beliefs are assertions. Perhaps the situation is different when we express a partial belief, but it's hard to see how. And it is hard to see how we could say that probability sentences express anything else

⁹ It should be remembered that Type 2 theories themselves form a large class, so that Lewis, Keynes and Carnap appear to all endorse theories from this class, but this does not imply there is close similarity between their respective views.

without giving up the hope of analysing probability sentences in terms of merely rectified beliefs, rather than say rational belief. So for these three reasons type 3 subjectivism seems untenable.

1.8.4 Type 4

Moving to type 4 subjectivism solves all three of these difficulties in one stroke, which bodes rather well for its success. The idea behind this theory is that probability sentences express commendation of certain epistemic states and disapproval of others, or perhaps express some more subtle dispositions to commend and disapprove. The idea is just to extend the analysis of ethical sentences offered in Ayer (1936), Blackburn (1984) and Gibbard (1990) to probability sentences. Indeed, Gibbard explicitly endorses an expressivist analysis of ‘reasonable’ and Blackburn (1980) suggests a very similar account of ‘chance’, though he uses the term in much the way that ‘probability’ would now be used.¹⁰ Blackburn claims that Ramsey also adopted this account, which might be correct.

One unimportant technical point before proceeding. We can’t analyse ‘The probability of p is 0.2’ as a commendation of believing p to degree 0.2. The simple reason is that ‘The probability of p is 0.2’ entails ‘The probability of p is not 0.3’ but we can commend believing p to degree 0.2 without disapproving of believing p to degree 0.3, because in some circumstances we might regard different, and indeed incompatible, states worthy of commendation. Similar remarks apply if we analyse probability sentences as expressions of something more complicated. But this problem is solved by just analysing an utterance of ‘The probability of p is 0.2’ as commendation of believing p to degree 0.2 and disapproval of all other degrees. Anti-expressivist, or cognitivist,

¹⁰ For example, he argues that we can talk about non-integer chances in a deterministic world. We can certainly talk about probabilities in a deterministic world, but standard usage now seems to be that we can’t talk about such chances. See Lewis (1986: 118).

analyses of probability in terms of reasonable beliefs will have to make a similar complication to their story, so this is no argument against expressivism generally.

That unimportant point aside, there is a more pressing difficulty for expressivist theories generally. I won't go into great detail here, in part because a fair discussion of this point would require a thesis length exposition on its own. In part, however, my lack of detail is caused by the possibility that there is less distance between my position and the expressivist position than appears at first. Some modern theorists, including some disposed to expressivism, have thought that an expressivist approach to some class of utterances, ethics being most frequently discussed, is compatible with believing utterances in that class to be truth-apt (e.g. Price (1994), Horwich (1994)). Since the traditional statement of expressivism is precisely that certain classes of utterances are not truth-apt, this might seem like a fairly substantial change, but there are reasons for the move¹¹. Price, for example, argues that the essence of expressivism in ethics lies in the claim that the function of moral utterances like 'Stealing is wrong' is significantly different from the function of non-moral subject-predicate sentences like 'Snow is white' despite their common syntactic form. The latter class have as their primary aim making (accurate) descriptions of the physical world; moral sentences have as their primary aim expressing a certain outlook. If Price is right then the difference between Type 2 and Type 4 theorists lies only in the *pragmatics* of probability sentences, not in their semantics, or for that matter their syntactical rules. This is undoubtedly an important question, but it's not one I've sought to address here. So I regard this type of expressivism as compatible with the theories I'm promoting.

¹¹ For Price, it is to escape from the Frege point I'll set out presently; for Horwich, it is because of his minimalist conception of truth.

The important problem for expressivism is what has become known as the Frege point. This was first explicitly set out in Geach (1965), though Geach had hinted at it earlier. The point is that the following argument is clearly valid.

- (1) If stealing is wrong, then getting little brother to steal is wrong.
 - (2) Stealing is wrong.
-
- (3) Getting little brother to steal is wrong.

There are two problems intertwined here for the expressivist. The first is explaining how we get the meaning of (1) from its components. That is, it clearly isn't a full explanation of the meaning of (2) to say that when uttered it expresses a con-attitude towards stealing, for this doesn't explain how it contributes to the meaning of (1). This is a decisive refutation of some primitive expressivist theories (like that in Ayer (1936)) but is no problem for modern approaches which acknowledge this question and present answers to it. However, it is a constraint on those answers that they be consonant, in some broad sense, with the expressivist analysis of (2). In part this consonance is imposed for theoretical considerations; it would hardly be plausible to say moral words like 'wrong' function in a radically different way in antecedents to the way they function in simple sentences. And it's imposed because of the second problem for the expressivist; they have to explain how the argument is valid. That is, they have to show the logical incoherence of accepting (1) and (2) and not accepting, or worse denying, (3). And in part this will require showing there is no equivocation in meaning between (1) and (2), else we will not have a clearly valid argument.

Price (1994) suggests we can get out of this with an expressivist analysis of conditionals. His theory of conditionals might be on the right track, and seems to do the work the expressivist needs¹². But I don't think this solves the overall problem. The problem arises because of a

¹² See Barker (1995) for an outline of a pragmatic theory of conditionals that seems broadly correct and

convergence of two facts: moral sentences occur in unasserted positions in sentences, and those sentences combine with simple moral sentences to form valid arguments. This is exemplified by conditionals like (1), but it is also exemplified by disjunctions like (1').

(1') Either stealing isn't wrong or getting little brother to steal is.

The truth functional analysis of conditionals has prominent supporters, but it is highly controversial and it hardly seems to be a refutation of a theory that it needs to deny it. On the other hand the truth functional analysis of disjunction is so entrenched, and so explanatorily successful, that it would require some large trade offs for it to be given up. And disjunctive syllogism is slightly more contentious than *modus ponens*, but still commonly enough accepted that it would be a cost for the expressivist to give it up. So I suspect the expressivist has to explain how moral sentences function as disjuncts, and how this story combines with the ordinary story about disjunction and validity to yield the validity of the argument from (1'), (2) to (3).

The problem has been the subject of a number of attempted solutions. However, I agree with Hale's contention that these solutions fall to a simple dilemma (Hale 1993: 340). Either the solutions do not explain how arguments like (1') and (2) to (3) are logically valid in the sense that a person asserting the premises and denying the conclusion would suffer from a *logical* shortcoming, or they fail to explain the meaning of the disjunction in a way consonant with the expressive explanation of the meaning of the disjuncts. Hale argues that the solution proposed by Blackburn in his (1984) falls to the first horn, and the new solution proposed in Blackburn (1988) falls to the second.

In (1984) Blackburn argued that we could interpret (1) (and (1')) as expressions of a con-attitude towards disendorsing stealing but not disendorsing 'getting little brother to steal'. The problem

compatible with Price's version of expressivism.

with this approach, as Blackburn came to realise, was that it posits the wrong kind of incoherence on the part of the person who asserts the premises and denies the conclusion. Such a person seems to suffer from the moral fault of not upholding their own second-order principles, but this is hardly a logical fault, which is what the expressivist needed to show. Blackburn has subsequently developed a different approach to explaining (1) and (1') (Blackburn 1988). Hale argues that the interpretation adopted there is ambiguous, either disjunctions and conditionals are read truth-functionally, in which case we don't have a reading consonant with the expressivist reading of simple sentences, or they are read expressively, in which case they still don't underlie the validity of the relevant arguments. There are more arguments to be had on this point, but there are enough problems here to suggest there is value in exploring a non-expressivist approach, as I do in subsequent chapters.

1.8.5 *The Exchangeability Point*

Given the objections I've made to subjectivism, the following defence of subjectivism may not seem immediately relevant, but perhaps its proponents intend it to defuse Ayer's objection that subjectivism doesn't allow for the obvious fact that epistemic states can be coherent but mistaken. In any case, the point may provide some defence of the Type 2 theory I want to defend. The idea is that coherence alone requires convergence of degrees of belief over time, so perhaps the epistemic states I described as coherent but mistaken are not really coherent at all.

The crucial concept is de Finetti's idea of *exchangeability*. I'll just deal with a very simple version of this idea, because it does well enough at bringing out all the philosophical points involved. Assume that m trials will be conducted, each trial having two possible results, say that for some variable x either $x = 0$ or $x = 1$. So there are 2^m possible outcomes for the series of trials. That is, we identify outcomes with the sequence of values of x according to each trial. Call the sum of an outcome the number of ones it contains. An agent regards the trials as exchangeable over this sequence of trials iff they have the same degree of belief in any two outcomes with the same sum,

and exchangeable generally iff they would regard any sequence of m trials as exchangeable, whatever the length of m .

Exchangeability is not the same thing as probabilistic independence¹³. An agent can regard the trials as highly interdependent in the sense that once they learn the outcome of an initial sequence of trials they would change their degrees of belief about the results of subsequent trials. For example, assume a biased coin is about to be tossed 5 times, with $x = 1$ meaning it lands heads and $x = 0$ meaning it lands tails. An agent regards the coin as so biased that she is certain it will land the same way on every trial. But she has no idea of the direction of the bias, so she assigns probability $1/2$ to the sequence $\langle 1, 1, 1, 1, 1 \rangle$ and $1/2$ to the sequence $\langle 0, 0, 0, 0, 0 \rangle$. Then she regards the trials as exchangeable in this sense, but clearly not independent.

The importance of exchangeability lies in some convergence results developed by de Finetti. Assume two agents update their beliefs by conditionalisation¹⁴, and regard a long sequence of trials as exchangeable. Then, provided they don't completely rule out some possibilities to start with, their degrees of belief about success on the next trial will converge. That is, for any $\varepsilon > 0$, there is an n such that after n trials their degrees of belief in success on the next trial will differ by at most ε . One philosophical interpretation is to say that this removes the more perniciously subjectivistic elements from subjectivism. The subjectivist now has an explanation of not just why convergence of opinion occurs (most dramatically perhaps in the convergence of opinion about decay times for radioactive elements), but of why it ought occur. Perhaps, the argument could continue, anyone who differed from this great convergence would be unreasonable in a way that even Type 1 subjectivists could object to.

¹³ Formally propositions A and B are probabilistically independent iff $Pr(A \& B) = Pr(A) \cdot Pr(B)$, or, equivalently, $Pr(A | B) = Pr(A)$.

¹⁴ That is, upon learning B they assign to $Pr(A)$ whatever value they used to assign to $Pr(A | B)$.

The problem with this move is simply that there is nothing in Type 1 subjectivism which grounds the claim that agents should regard certain trials as exchangeable. Indeed, in seeking to discriminate between different coherent states on the grounds of their reasonableness (i.e. between those that do and don't regard trials as exchangeable) we have slipped towards Type 2 theory, which Carnap showed is not subjectivist at all. This point is made by Kyburg (1978: 67) who attributes it to discussions with Nagel.

Matters are even worse for the Type 1 subjectivist. An event can be interpreted as many different types of trial. For example, drawing an emerald from an urn can be regarded as a trial of whether the emerald is green or not-green, and whether it's round or not-round. More interestingly, we can regard it as a trial of whether the emerald is grue or not-grue. And of course once we've recognised grue we can recognise all sorts of other predicates, such as green on an even numbered trial or blue on an odd numbered trial. We can't, consistently, regard all such sequences of trials as exchangeable. So we must make a selection, before the evidence comes in, as to what we will regard as the exchangeable trials. But that we must make such choices before seeing any evidence is what distinguishes Carnap's Type 2 approach from de Finetti's Type 1 approach.

Indeed, we can turn around de Finetti's result to be a defence of a variant of Carnap's position. The Type 2 theorist is burdened by the necessity of saying something about what is reasonable on zero evidence. Carnap rose to that challenge by trying to give the precise numerical value of every proposition on zero evidence, but as we saw in section 1.7, his attempts seemed doomed. Keynes allowed more flexibility by letting probability values be non-numerical, and I'll essentially be following Keynes here. I think Carnap's program is best served by not trying to find *the* reasonable probability function, but the set of such reasonable functions. de Finetti's convergence theorem can be used to argue that what distinguishes elements of this set is not the value they give to particular propositions under no evidence, as Carnap thought, but what sequences of trials they regard as exchangeable. Roughly, reasonable probability functions are

reasonable by virtue of their content, not as Carnap thought by virtue of their form. We are, however, getting ahead of ourselves. I'll return to this matter in my defence of this theory against various objections in chapter 6.

Chapter 2

What Degrees of Belief Aren't

Ramsey objected to Keynes's view that *probability* means reasonable degree of belief on the grounds that Keynes had provided no explanation of what degrees of belief were. So he proceeded to provide such an explanation, what I'll call the betting analysis. The point of this chapter is to show that Ramsey's analysis can't work, and that even if it did work the Dutch Book arguments based upon it are unsound. With this a major motivation for the betting analysis disappears.

2.1 Analyses of Degree of Belief

At his first, and most famous, attempt Ramsey said that having degree of belief r in A is thinking the bet $(1 - r, A, r)$ ¹ is fair. That is, its expected worth is zero, and an agent who makes this evaluation will be prepared to buy this bet for any price less than zero, or sell it for any price greater than zero. I'll call this the betting analysis of degrees of belief. There is an important qualification to this analysis. I haven't specified in what units the payouts are quantified. If the payouts are quantified in units like dollars with a declining marginal utility, what I'm calling Ramsey's approach won't work. It only works when the payouts are quantified in 'utils' or something equivalent. We can get around this problem in two ways. First, we can follow Ramsey and work out the utility of every possible outcome. Alternatively, we could set the payouts in a good which ought have constant marginal utility. Savage (1954) and Smith (1961) suggest that lottery tickets have this property, and hence develop their theories using lottery tickets as a currency.

There are, however, approaches to defining degrees of belief other than in terms of evaluations of bets and dispositions to bet. Indeed, one of these is set out by Ramsey himself in a later paper. For this approach we have to assume that people can compare, introspectively, their degrees of

¹ The bet (x, p, y) is the bet which pays x if p and costs y otherwise.

belief in different propositions. In his earlier paper Ramsey puts forward some arguments against this, but I think these arguments have to fail. To see why we just have to consider Ramsey's arguments carefully.

Ramsey discusses and rejects various possible introspective feelings that could serve as degrees of belief. One of these is 'intensity of feeling' (1926: 169). Ramsey simply rejects this by an example. We have much more intense feelings about some beliefs which, if pressed, we would admit we believe to a lower degree than those things which we take for granted. For example, in terms of intensity, my belief that 'Lowering tariffs improves welfare' is stronger than my belief 'The earth is round'. On the other hand, my degree of belief in the latter is higher than in the former. Indeed, I suspect my degree of belief that the latter is true is stronger than is my degree of belief that the former is even truth-apt. So my degrees of belief are not mapped by my intensity of feelings. I suspect that for most people we can find examples showing the same effect.

So this particular argument of Ramsey's is effective. But look at what we take as evidence. I simply accepted the introspective evidence that my degree of belief in those things which I take for granted, such as 'The earth is round' is high. Now perhaps even if we didn't have such introspective evidence we could still run Ramsey's argument by looking at the external evidence, such as betting behaviour, to determine our relative degrees of belief in the two propositions. However, I suspect no reader actually did that or anything like it when considering the examples. As Mellor (1980) notes, this point of Ramsey's serves to highlight the pretheoretic plausibility of saying we can introspectively make *qualitative* judgements about our degrees of belief.

This is an important positive argument for the analysis in the next chapter, but it is also a negative argument against the betting analysis. Because Ramsey thinks degrees of belief don't relate to any property determinable by introspection, he thinks that we have to look at what causal impact they have. That is, we have to look at their behavioural implications. Given this the betting

analysis seems the most plausible candidate. I agree that if “the kind of measurement of belief with which probability is concerned is ... belief *qua* basis of action” (1926, 171) the betting analysis seems the only plausible approach. However, the above argument, and the existence of the analysis developed in the next chapter, make this premise dubious. The reason the existence of a positive analysis is important is that the early Ramsey denies that there is any ‘introspected feeling’ which could be measured in the right type of units to be probability. The equivalence analysis shows that we can find such a feeling.

2.2 Why the Betting Analysis of Degrees of Belief Fails

Since the definition of degrees of belief in terms of bets is the orthodoxy, I have to show why I think it is untenable before I can justify an alternative outlook. Before going into detail about why I think the identification of degrees of belief with propensities to bet is wrong, I’ll simply list the objections:

- Our propensity to bet is affected by our attitudes towards gambling.
- Our intuitions about what is reasonable to do on the assumption that the marginal utility of the currency is constant are distorted by our intuitions about everyday situations.
- On a betting analysis we can’t get our probability logic in order until we have worked out our logic of preference, but this has many foundational difficulties.
- In betting situations there is no operational difference between a proposition being true and its truth being discovered, and if these are significantly different, the betting mechanism will determine our degree of belief in the wrong proposition.
- The betting analysis presupposes that norms of practice are epistemic norms, but if this is true it is something which must be proven not presupposed.
- For at least *some* incoherent agents, the analysis looks like it gives the wrong answers.

- The betting analysis cannot support Dutch Book arguments, as is commonly assumed, and hence does not ground the norms we want, whereas the analysis in this dissertation does ground these norms.

The first objection is a common one, and perhaps not too serious. It has been suggested that we get around this problem by not looking at whether people would bet but rather at which gambles they would accept if offered a choice from a set of desirable gambles². The latter approach runs into problems of its own because, in cases where agents are susceptible to Dutch Books, the problem won't be the relatively serious one that they can be made to lose money, but the relatively trivial one that they will receive a smaller gift than they could have. Even in these cases, the experimental evidence is that people have a preference for choices which seem to involve less of a gamble, in some undefined and possibly incoherent sense³. The worry is just that these dispositions for or against gambling itself will pollute our information about degrees of belief. Indeed the fact that we can, it appears, sensibly talk about attitudes to gambling polluting the information about degrees of belief seems to count against the idea that we can *analyse* degrees

² This move is made in Savage (1954: 28). He credits de Finetti (1937) as the inspiration for it.

³ The experiments I am thinking of are of the following form. Assume we have two goods, *B* and *C*, such that *C* is considerably better than *B* but not overwhelmingly so. If we aren't paying our experimental subjects, *B* could be \$1 million and *C* \$5 million. If offered the choice between a 20% chance of *C* and a 25% chance of *B*, subjects will, on the whole, choose the former. The same choice can be seemingly set up as a two-stage process. Whatever the subjects choose, there is a 25% chance of qualifying for the 'second round'. If they qualify they will receive either *B* for certain or an 80% chance of *C*, but they have to say before they know whether they've qualified which they would choose. Even though this seems functionally equivalent to the first choice, here subjects will overwhelmingly choose *B*, because it doesn't involve a gamble. There is no consequentialist sense in which these choices are coherent. These experiments are reported in Kahnemann and Tversky (1979: 273ff). They refer to the choice pattern exhibited as 'the isolation effect'. I'll discuss the relationship between dynamic and static choices in appendix 3C and also in chapter 9.

of belief as dispositions to gamble. Certainly attitudes to gambling do not pollute any information we might get about dispositions to gamble. Perhaps though this last point merely shows that the betting analysis is not obviously true, not that it is false.

The second objection is related to the first, but it is more pragmatic. Ramsey argues that what we mean when we say a belief is reasonable is that it was formed by a reasonable habit (1926b: 194ff). But what habits are reasonable depends in part on our surroundings. In particular, our reasonable habits may become unreasonable if we are placed in a radically different situation. A similar situation arises in ethics where if we think behaving ethically is behaving in accord with certain sets of rules we have to acknowledge the possibility that in certain unusual situations good actions will have less desirable expected outcomes than bad actions. The objection is that in situations where we are making bets in ‘utils’ are so different to everyday life that what is reasonable in our situation may be unreasonable in those. Hence our intuitions about what would be reasonable seem unreliable. Since the main point of the betting analysis is to work out what is reasonable, and our basic data is our intuitions about the reasonableness of specific acts, this vitiates the usefulness of the analysis.

On a betting analysis, we work out what degrees of belief are reasonable by looking at what preferences are reasonable. This, however, increases unduly our workload in the foundations of probability. For example, Good (1952) notes that we have to work out how to deal with infinite utilities (or show their impossibility) as a foundational task in probability logic on a betting analysis. I don’t doubt there are various plausible ways of doing this, but I would prefer my theory of probability was *not* held hostage to a particular analysis of infinity if possible. For a different example, Savage’s axioms of preference in his (1954) are much more contentious than the probability logic he derived from them. The point is simply the pragmatic one that the less contentious philosophy we have in our foundations the better.

When we place a bet, we don't care directly about whether or not the proposition on which we bet is true. On the contrary, we care directly about whether we and our bettor will come to know that it is true, so we can claim our winnings. So the betting analysis of degrees of belief should lead to our logic of degrees of belief being intuitionist not classical (Harman 1983). Ramsey gets around this by assuming the bookmaker has the power of the Almighty. It would be distressing if our commitment to classical logic depended on being theists. (There is an interesting comparison here with Dummett's arguments in 'Truth' for the claim that opposing intuitionism requires some kind of commitment to the supernatural.) Even if we accept that it is possible for the bookmaker to have this power, the concern from the last paragraph about the usefulness of our intuitions in these circumstances remains.

It has been pointed out by several authors that Dutch Book arguments grounded on the betting analysis, even if sound, make a large presupposition⁴. That presupposition is that norms of action, such as 'don't buy Dutch Books', are epistemic norms. Now it may be possible to prove that prudential norms are epistemic; indeed I suspect the analysis here goes some way to proving that. But it is a surprising result, and one for which we ought develop arguments. It isn't something that can be safely supposed, as it appears to be under a betting analysis.

Christensen (1996) argues for what he calls a 'metaphysical separation' between degrees of belief and betting practices. At one level his argument is an old-fashioned open question argument. Even if we know that someone has degree of belief 0.2 in p , we can still ask what evaluations they would make of bets on p . Now such arguments aren't particularly telling; it's no refutation of an analysis that it isn't obvious. Christensen, however, has in reserve a somewhat more subtle argument. Assume I will pay 30 cents for the bet $(\$1, p, 0)$ but only 20 cents for the bet $(\$1, p \vee q, 0)$. My evaluations are, in a sense I will get to, incoherent. The best justification for the

⁴ For a recent example, see Kvanvig (1994).

betting analysis is that we have to identify mental states by their functional role. If we were to assume the only role a degree of belief plays is in bet-evaluation, there might be a functionalist argument for the betting analysis. As Christensen notes, however, in this case my degree of belief in p , whatever it is, performs at least two roles. One is in helping determine how much I'll pay for $(\$1, p, 0)$, and the other in helping determine how much I'll pay for $(\$1, p \vee q, 0)$. In fact there is a third role; helping determine what my degree of belief in $p \vee q$ is. When we're coherent, there will be no tension between these roles. But incoherence, at least of preference, is clearly possible. Christensen's point, I take it, is that for incoherent cases the betting analysis unjustifiably privileges one particular functional role to the exclusion of others, and hence it can be challenged on its own ground.

Finally, one of the advertised strengths of the betting analysis is that through Dutch Book arguments we can provide a justification for degrees of belief obeying axioms of the probability calculus. This was originally argued by Ramsey, and has been extended to dynamic settings by Lewis and van Fraassen. If these arguments succeed they show that the betting analysis has great practical usefulness. However, I show in the next section that these arguments are invalid, or at least unsound. On the other hand, I prove in the next chapter that all the results which Dutch Book arguments claim to achieve can be grounded in a purely epistemic analysis.

2.3 Dutch Book Arguments Fail

For simplicity, let's define *the A-bet* to be the bet $(1, A, 0)$, where the unit is of some currency with constant marginal utility and $Bel(A)$ to be my degree of belief in A . The following is a paradigm Dutch Book argument. Assume $Bel(p)$ is 0.6 and $Bel(\neg p)$ is 0.55. Then, by the betting analysis, I will be prepared to pay $Bel(A)$ units for an A bet, or at least $Bel(A) - \varepsilon$, for arbitrarily small ε . We'll assume for the sake of the argument that the marginal utility of money is constant

in small amounts⁵. Hence I will pay \$1.15 for a p -bet together with a $\neg p$ -bet. The sum of these bets is $((\$2, \neg p, \$1), p, (\$1, \neg p, 0))$ or, in other words, a bet which pays \$2 if p and $\neg p$, \$1 if p or $\neg p$ but not both, and nothing if neither p nor $\neg p$. Since it is impossible that p and $\neg p$ it is impossible that this bet will pay more than \$1, hence my betting practices are normatively flawed. And since by assumption norms of betting behaviour are epistemic norms, I must be irrational. Such arguments can be used to show that my beliefs ought obey all the axioms of the probability calculus.

I agree that if I was prepared to pay 60 cents for a p -bet and 55 cents for a $\neg p$ -bet this would be irrational. That is, I accept for the sake of argument the identification of norms of betting with epistemic norms. However, I don't see how it follows from having certain degrees of belief that I ought be prepared to pay these amounts. The problem is that the argument assumes I will not use any strategic pricing, yet it gives no reason for thinking that I oughtn't price strategically. Indeed it seems implicit in the argument that I ought price strategically. By strategic pricing I mean setting a price for a good that is not determined solely by its intrinsic usefulness, but by how much I could either sell the good for or obtain the good from other sources.

Adam Smith noted some 220 years ago that the price of goods bore no interesting relationship to their usefulness. Nothing is more useful than water, yet it is almost free, nor less useful than diamonds, but they have massive value. We have since learnt that there is in some specified

⁵ This is a common enough assumption in this field, but I don't see any economic reason for it. If we assume, as is standard, that the utility of any particular level of wealth is independent of our actual wealth, there is no reason to think that the marginal utility of money will be constant around our actual position than it would be around any other positions. If, on the other hand, the marginal utility of money in small quantities really is constant, this leads to problems for orthodox utility theory. I don't want to argue for either of these assumptions and against the other, but it is worth noting that despite the frequency with which each is assumed they are in strong tension.

circumstances a determinate relationship between what we'll call the value of a good and its usefulness. If a consumer's budget is at equilibrium then the ratio of the marginal utility of a good to its marginal price will be constant for all goods the consumer purchases provided the utility function of every good is differentiable (Slutsky 1915). Still, in general, i.e. at disequilibrium, Smith's observation holds. Moreover, since we are assuming that, for the bets in question, the marginal utility of both the payouts of the bets and the currency we use to buy them is constant, unless all people have the same degree of belief in all propositions we can't trade our way to an equilibrium position.

The principle flaw in Dutch Book arguments is that they ignore Smith's observation. My degree of belief in A can determine at most the usefulness of an A -bet. Yet it is assumed that it will also determine the price I am prepared to pay for A -bets. Hence it is assumed that there is a correlation between how useful bets are and how much I will or ought be prepared to pay for them. As Smith showed, in general this cannot be the case.

We can find simple examples where it would be unreasonable to pay $Bel(A)$ for an A -bet. Assume that I know I can sell a p -bet for 90 cents, and $Bel(p)$ is 0.7. I am offered a p -bet for 80 cents, should I accept? According to the betting analysis I should not, because I am being asked to pay 80 cents for something which has an expected value of 70. However, it seems at least plausible that I should accept the bet and then sell it for a sure profit. Alternatively, assume I know I can buy as many p -bets as I like for 70 cents each in the market, and $Bel(p)$ is 0.9. Again I am offered a p -bet for 80 cents. The betting analysis says I should accept, but again it's plausible that I should instead buy p -bets at the cheaper market price.

This may not look at first like a major difficulty. After all, we know the correlation between fair prices for A -bets and degrees of belief in A only holds under restricted conditions. All these examples show is that we have to be more careful in specifying the initial conditions. As far as it

goes, this response is correct. Provided we have good reason to believe that we oughtn't use strategic pricing, the correlation will hold. The problem for Dutch Book arguments is that the only way we can know this is if there is no possibility of later bets, so we can know that we can't buy the bets for less on the market nor make a profit by resale. However, as we saw above Dutch Book arguments in general only work by using retrade. Hence they rely on the correlation between degree of beliefs and betting prices for rational agents holding in a context in which only irrational agents would price bets this way, so the arguments fail.

This conclusion is of major importance for what follows, so I should restate the argument which I have used. Dutch Book arguments rely on there being entailments from agents having certain degrees of belief to their propensity to buy certain bets. They conclude that if some set of our degrees of belief are probabilistically incoherent, we will buy a set of bets which incurs sure loss, and hence we must be irrational. However, the entailment in question only holds under restricted circumstances. One of the restrictions is that there be no possibility for later trade in bets. When there are retrade possibilities, as there must be for most types of Dutch Books to be made, the entailment does not hold. So Dutch Book arguments make inconsistent presuppositions, and hence fail.

It might be objected that if an agent whose beliefs were not coherent with the probability calculus believed falsely that there was no possibility of retrade they would make trades which led to sure loss. However, the only way a bookmaker could exploit this is if she had more knowledge than the agent. And the fact that a bookmaker with more knowledge than us can sell us bets which, given the bookmaker's knowledge, have to lose, is no proof that we are irrational. If it was we would be able to show that any person whose degree of belief was less than 1 for any true proposition, or greater than 0 for any false proposition, is irrational. Moreover, even if we regard such an agent as irrational, it is not clear that it is because her beliefs don't follow the probability calculus that she's irrational. All we can tell from the fact that she will suffer a sure loss is that

she's made a mistake somewhere; this might be concerning her misplaced certainty that the market is closed rather than her degrees of belief in the proposition on which bets are placed.

Alternatively, it might be objected that an agent who is completely ignorant of the state of the market will price bets by their expected return even if they think there is the possibility of retrade. The problem with this objection is that it is, famously, very hard to pin down what it is to be completely ignorant. Saying that if I am completely ignorant of whether or not it is the case that p then my degree of belief in p is, or ought to be, $1/2$ leads to well-known contradictions. It certainly would be odd to say that if we are completely ignorant of the likely effects of a certain class of events we should ignore them, which would seem to be the line of attack here. In part 2, particularly in chapter 9, I'll look at various theories about how we ought make decisions under ignorance. Under several of these (particularly maximin approaches) if the agent knows nothing about the market she should make no trades at all rather than pricing according to mathematical expectation. Finally, if the approach advocated in this dissertation is correct and there are necessary probabilities, then in many circumstances an agent is irrational to be completely ignorant in some strong sense. So again, even if the agent is irrational, the Dutch Book argument can't prove that it is the incoherence of degrees of belief with the probability calculus that is making the agent irrational.

In the literature there is generally a distinction drawn between synchronic and diachronic Dutch Book arguments, with the latter being referred to as Dutch Strategy arguments. These latter type are used to infer coherence constraints on how our degrees of belief should change over time. Though the above argument refutes both Dutch Book and Dutch Strategy arguments, the results concerning Dutch Strategy arguments are more striking. Unlike Dutch Book arguments, Dutch Strategy arguments appear to have occasionally led to authors drawing mistaken conclusions.

In his latest argument for a principle called Reflection, van Fraassen discusses the case of Pierino, whom he claims is irrational (1995: 11). Pierino is a young child who today prefers blocks to marbles, but knows that in a year when he has acquired older tastes, he will prefer marbles to blocks. He is in the predicament of today having 9 marbles. Van Fraassen stipulates that Pierino is indifferent between keeping his 9 marbles and trading them for 3 blocks, knowing full well that in a year's time he'll be indifferent between holding these blocks and trading them for a single marble. And this, van Fraassen claims, must be irrational, for if he made the two trades he would have lost 8 marbles. To the obvious response that he will have gained in enjoyment in the short term by having more blocks which he can use now, van Fraassen replies that since he was indifferent to the trades, he can't gain anything.

van Fraassen's reply makes the mistake we have attempted to highlight here. Pierino's indifference to the trades tells us that at each time the exchange-value of the bundles on offer was equal, but we can't from this infer that the use-value of the two bundles was equivalent. Indeed, assuming Pierino preferred more marbles to fewer in year 2, we must assume that the value of having 3 blocks in year 1 was greater than the value of having 9 marbles, and indeed so much greater that it made up for the expected losses in year 2.

We can make a formal model for Pierino that satisfies these constraints. Assume that U_1 is the utility he gets from toys in year 1, and U_2 the utility he gets in year 2. Assume his aim is to maximise U_1U_2 , and hence that he is indifferent as to the amount of toys he holds at the end of year 2. Let B_i and M_i be the amounts of marbles and blocks he has in year i , and assume his utility functions are as follows.

$$U_1 = 27B_1 + M_1;$$

$$U_2 = B_2 + 3M_2.$$

Given that he starts with 9 marbles, if he just holds marbles his net utility across the two years will be $9 \cdot 27 = 243$. If he trades the 9 marbles for 3 blocks, knowing that these can only be traded for 1 marble at the end of year 1, his net utility will be $81 \cdot 3 = 243$. Hence his indifference to the trade. He could increase his utility if it is possible to trade say 6 marbles for 2 blocks, but we have no reason to assume that that trade is allowed. For reasons I will outline in the next chapter, I think van Fraassen's main conclusion, that all rational agents are Reflective, is sound, but the arguments he uses to get there are mistaken.

One final point ought be noted. When an agent holds a Dutch Book what is important is not that there is no winning outcome. What's important is that all the winning outcomes are impossible. Assume as above that I bought a p -bet for 60 cents and a $\neg p$ -bet for 55 cents. Then if it's the case that p and $\neg p$ I will win 85 cents, and if it's the case that neither p nor $\neg p$ I will lose \$1.15. However, in all possible situations I will lose 15 cents. This has two interesting consequences.

First, Dutch Book arguments presuppose a certain logic; in particular one where p & $\neg p$ is impossible and $p \vee \neg p$ is a tautology. There's nothing wrong with this, but it is a presupposition which should be noted. (Harman (1983) notes that because of this attempts to use the probability calculus to prove the semantics for natural language ought be classical are question-begging). Unless otherwise stated, I will presuppose classical logic. That is, when I say A is impossible I will mean $\neg A$ is a classical theorem, and when I say A entails B I will mean it classically entails B .

Secondly, there is a *sound* Dutch Book argument which can be made against a person whose degree of belief in any contradiction is positive, or whose degree of belief in any tautology is less than 1. I'll just illustrate the first. Assume my degree of belief that the four-colour map theorem is false is 0.1. Then I'll buy a bet against it for 5 cents. As there's no possibility of this bet winning, this is a sure loser. And since there's only one bet involved, I didn't assume that retrade was

possible, so the above objections to Dutch Book arguments don't apply. The objector who says that it is an *epistemic* possibility that the four colour map theorem is false has a much deeper objection to Dutch Book arguments than I. After all, p & $\neg p$ might be an epistemic possibility too, so on this approach we could object directly to the toy Dutch Book argument with which I opened this section. Rather than take this road, I will simply assume that in this field we are interested in an epistemology for agents whose beliefs are closed under entailment.

2.4 Other Critiques of Dutch Book Arguments

The critique of Dutch Book arguments given here is unique in two respects. First, it is the only one, to my knowledge, to rely on Adam Smith's distinction between usefulness and exchange-value. Secondly, as will be seen in the next chapter I concur with the most famous conclusions of Dutch Book arguments. The usual motivation for criticising these arguments is to motivate dissent with their conclusions. Here the motivation is to provide a cleaner separation of epistemology and decision-theory.

The closest argument in the literature to mine is given by Schick (1986). He argues that Dutch Book arguments fail because they assume that the value of bets is additive. That is, they assume the value of an A -bet is independent of whether or not the agent holds a B -bet. Since bets might be complementary in an economic sense, this is a false assumption, so these arguments fail. This objection is similar to mine in that it relies on a simple economic theory to refute the Dutch Book argument, and because as Schick notes it doesn't apply to 'single-bet' books (1986: 116). However, it is not a successful refutation.

The alleged flaw with Dutch Book arguments on which Schick relies was noted by Ramsey when he originally put the argument. (Ramsey 1926: 173-4). Not only does Ramsey point out the alleged flaw, he notes its *prima facie* implausibility and offers a small argument to try and defend

it. Ramsey claims that when all the final payouts⁶ of all bets are ‘ultimate goods’, the value of the bets is additive. Now Ramsey’s claim here might be wrong, but we should get an argument to this end. Instead, Schick simply assumes that when bets are denominated in utils (equivalently when the marginal utility of money is assumed to be constant) we will get similar economic results to those we’d get were bets denominated in dollars. Schick’s mistake (if it is a mistake) is instructive; as I noted above the fact that we don’t, even on reflection, have particularly clear intuitions about trading utils is a good reason for not founding our theory of probability on the types of bets Ramsey discusses.

Levi (1987) and Maher (1992) argue that agents who are not Reflective will ‘see the Dutch Book coming’ and hence refuse to take the bets which lead to being Dutch Booked. Their argument is principally developed to defeat van Fraassen’s conclusion that ideally rational agents are Reflective, though it isn’t clear why it wouldn’t also apply to synchronic Dutch Book arguments. Since it doesn’t actually work it isn’t particularly worthwhile to speculate how far it would reach were it successful.

The kinds of cases they are thinking of are like the following. Assume I today believe that the probability of p is 0.5, and believe that tomorrow I’ll believe the chance is 0.3. Assume also I’m offered a p -bet for 40 cents. I know that if I buy it I will be prepared to sell it tomorrow for say 32 cents, for a sure 8-cent loss. That is, I’ll be Dutch Booked. But wait! If I see this coming I won’t buy the original bet for 40 cents, and thus avoid holding the book. Levi and Maher claim that the availability of this path to unreflective agents blocks the Dutch Book argument.

⁶ When I say the ‘final payouts’ of bets are of type X I mean the following. The set of all bets whose ‘final payouts’ are of type X is the smallest set of bets including all those whose payouts are of type X such that any bet such that each of the payouts is a bet in the set is also in the set. If we allow bets to have more than two possible payouts we can amend this last condition accordingly.

This response to Reflection fails for the simple reason that being ‘Dutch Bookable’ is not a necessary condition of irrationality. Assume I don’t buy the original 40 cent bet. I won’t now be able to sell this bet for 32 cents tomorrow. However, I will still be able to buy a $\neg p$ bet for 68 cents. If I take Levi and Maher’s advice, I’ll have converted a sure loss of 8 cents into an expected loss of 18 cents. There might be an argument to show that this is a rational option, but I’d like to see what it is⁷.

There is a bigger problem for Levi and Maher’s approach. As we saw above, when used by de Finetti and Savage the Dutch Book argument does not require the agent to incur an actual dollar loss. Rather, since the choices are between gifts, the incoherent agent incurs a sure *opportunity* loss. Now when I refuse the original offer of a p -bet for 40 cents, I have already incurred an opportunity loss. Admittedly it is again an *expected* opportunity loss rather than a sure one, but it isn’t clear why incurring an expected loss rather than a sure one is an epistemic improvement.

Bacchus, Kyburg and Thalos (1990) run a series of responses to Dutch Book arguments. Their responses to dynamic Dutch Book arguments will be discussed in chapter 3 as possible objections to my arguments for dynamic coherence; here I’ll stick to discussing their general comments on

⁷ The practicalities of this situation are very difficult, and it is impossible to get clear intuitions about what we should do given the assumption of constant marginal utility of money. If I know that tomorrow my degree of belief in p will be 0.3, ideally I will take steps to prevent myself acting on this later belief. That is, I should take the Ulysses option. Now buying p -bets today for 40 cents to sell tomorrow at a sure loss will, as is noted in the text, reduce my *expected* loss. However, that is assuming that the number of $\neg p$ -bets I will buy (number of p -bets I will sell) tomorrow is independent of the number of p -bets I buy today. The idea is that taking Levi and Maher’s advice *may* in some circumstances, have the effect of tying me to the mast and not trading tomorrow. It *will* have this consequence if the marginal utility of money is not constant, but when it is there are few clear intuitions on the matter.

Dutch Book arguments. Put in slogan form, they endorse the position that bad betting is bad betting, not bad believing. I agree, but I'm a bit worried about the way they get to this slogan.

Having incoherent degrees of belief (and even the disposition to convert these directly into bets) does not guarantee sure loss. Only this combined with a rather clever and devious bookie does. Note two important consequences of this qualification. First, we now say that certain sets of degrees of belief will not *always* lead to losses, but will sometimes lead to losses. But we knew all along that any degrees of belief (except certain kinds of dogmatic acceptance of only tautologies) *might* lead to losses. Why, we can ask, are the losses caused by devious bookies signs of irrationality, but not the losses caused by taking attractive but ultimately losing bets? I suspect this raises problems for a certain type of pragmatist, but I can't see it as a general problem. The problem isn't that some possibility claim, i.e. we might lose if a certain type of bookie exists, is true, but rather an existence claim, i.e. that a certain type of acceptable but losing bet exists. The second consequence Bacchus, Kyburg and Thalos draw is that the Dutch Book argument only works if we make the paranoid assumption that devious bookies exist. Consistency isn't just the sign of a small mind, but of a paranoid one too. Again, this looks like a good refutation of a certain strictly pragmatic Dutch Book argument. However, we don't need to formulate Dutch Book arguments as strictly pragmatic, and when we don't I suspect this objection loses its force. That is, the possibility of the agent buying a Dutch Book seems at least as great an epistemic flaw as actually making the purchase, and hence anyone who runs a Dutch Book argument is just making an avoidable mistake if they assume an actual pernicious bookie.

2.5 The Betting Analysis as Analogy

Although I don't regard the betting analysis as correct, or even that useful generally given the failure of Dutch Book arguments, it may be a helpful analogy. Much of this section is motivated by Shafer (1981), who also regards betting prices as an occasionally useful analogy to degrees of belief, though he's considerably more sceptical than I about the applicability of this analogy. To

see when this analogy might be useful, we first have to consider the known limitations on its applicability.

The only way I will use betting examples is to test whether $Bel(p)$ is reasonable by considering whether it is reasonable to accept or reject offers to buy or sell p -bets. If, for example, in appropriate circumstances it would be unreasonable to reject an offer of a p -bet for 0.2, this can be taken as a good argument for saying that $Bel(p) \leq 0.2$ is unreasonable. First we must consider what circumstances are ‘appropriate’, or since this seems a bit open-ended, which circumstances are known to be inappropriate. The following have already been mentioned.

- When the units in which the bet is denominated or traded are of variable marginal utility.
- When there is a time-delay between when the bet is or would be traded and when winnings would be paid.
- When there is a possibility of trading in other bets at a later time.

The last is actually a bit broader than what we used above. We have to rule out not just trade in this particular bet, but in other bets as well because some bets are complementary in the economic sense. This is very common in real life. For example, it is worthwhile to buy insurance on your car but not on someone else’s despite the fact that the cost of the bets are the same and the expected returns may well be identical (unless say you know you are a worse driver than other people). In part this will be because insurance bets are denominated in a currency with declining marginal utility. However, it seems presumptive to think that complementation is only caused by this effect. Experimental evidence suggests that for many agents the fact that their degree of belief in some propositions is vague leads to a complementation effect. So we have in general to assume this is the last possible trade.

This assumption also gets us around a problem noted by Pargetter and Davidson (1985). Assume I know that your degree of belief in p is 0.7, and mine is 0.9. You offer to sell me a p -bet for 0.85.

Assuming all other circumstances are in order, this trade will be worthwhile for me. However, I know that if I counter-offer to buy it for 0.71, you will still find the trade worthwhile, and I will have bought the bet for 0.14 less. Pargetter and Davidson thought we could only get around this by assuming the agent under investigation knows the bookmaker has the same degrees of belief as they. However, once we know this is the last chance to bet, i.e. that the counter-offer possibility is closed, we don't need to make this extra restriction, and since we need to have a closed market after the bet in question for other reasons, Pargetter and Davidson's restriction seems redundant.

Even though it is not needed in these cases, however, we might want to restrict attention to cases where each party to the bet is known to have the same information for the following reason. Degrees of belief can at most determine dispositions to bet, not actual betting practices. Even if I have a disposition to buy p -bets for 0.2, if I am offered a p -bet for that price I might not buy. This seems contradictory, but it is not. Dispositions can be finkish (Martin 1994, Lewis 1997). I might have a disposition to ϕ in circumstances C , yet be in a situation such that whenever circumstances C arise I will lose the disposition.

Assume I have next to no evidence about the players in a certain tennis match, and let p be the proposition that the player who serves first will win. Even if I have a disposition to buy p -bets for 0.1, say, if someone were to offer me a p -bet for that price I would most likely refuse. That is, the disposition would be finkish. The reason I would refuse is that the fact I was offered the bet would count as a new piece of information (the information that someone who most likely knows more about the match than I thinks p -bets are worth less than 0.1) and in the state with this extra information I'm not disposed to make the purchase. If I did have the disposition this would be just like paradigm cases of finkish dispositions because the occurrence of the circumstances which are meant to 'trigger' the disposition causes me to lose that disposition in an easily identifiable way.

How could we tell that I originally had a finkish disposition rather than having no disposition at all? The best test seems to be whether I would buy the bet if I knew that the person offering it had the same information I did, and hence that there was little information in the fact that the bet was offered. So the Pargetter and Davidson restriction to circumstances where the offeree knows the offerer has the same beliefs they do is important for cases of complete ignorance to eliminate the effect of finkish dispositions. This will be used in later chapters.

Given these restrictions, it seems the analogy with bets is worthwhile. We will have to be careful to use it only in appropriate circumstances, and to remember that it is only an analogy, and perhaps not the only one. Where possible it will be preferable to use the analysis of degrees of belief to be developed in chapter 3.

Chapter 3

What Degrees of Belief Are

3.1 The Equivalence Analysis

In chapter 2 I noted that I would take qualitative data about an agent's degrees of belief as given. So I have as data to work with sentences of the form *The agent's degree of belief in A is higher than in B*. Standardly in the literature, the agent is called You, with the capitalisation meant to indicate that *You* is being used as a name. Hence the possessive form of You, as used, is You's, but it is standard to use Your. Rather than mimic this faulty grammar I'll use a variety of agents as appropriate.

In his 1926 paper, Ramsey argued that there was no way to convert qualitative judgements of greater or smaller degrees of belief into quantitative judgements of, say, degree of belief $2/3$. However, there is a relatively simple way to do this, as he pointed out in a note in 1929. To say my degree of belief in A is $2/3$ is to say I have the same degree of belief in it as I have in $p_1 \vee p_2$ when I know that exactly one of p_1 , p_2 and p_3 is true and each of the p_i are equally likely. From the qualitative judgements that A has the same degree of belief as $p_1 \vee p_2$ and that all p_i have the same degree of belief, we can work out a quantitative judgement.

It is rather trivial to generalise this. My degree of belief in A is x/y when it is the same as my degree of belief in $p_1 \vee \dots \vee p_x$ given that I believe fully that exactly one of p_1, \dots, p_y is true and each p_i is equally likely. By equally likely, I just mean that my degree of belief in p_i equals my degree of belief in p_j for all i, j . We will look in section 3.10 at how this analysis might be extended to real-valued degrees of belief, but until then we'll assume that degrees of belief take rational-values only. (Note that in this chapter 'real' and 'rational' always refer to properties of numbers. Saying that a degree of belief is real doesn't entail that anyone has it, nor does saying it is rational entail that anyone should have it.) I'll call this analysis of degrees of belief, that to believe A to a certain degree is to believe it to the same degree as a certain disjunction, the *Equivalence Analysis*.

It might seem that the Equivalence Analysis is circular, since we have analysed degrees of belief in terms of, *inter alia*, degrees of belief. This objection misses the point somewhat. The aim of the Equivalence Analysis is to reduce quantitative degrees of belief to qualitative degrees of belief. If we were to write it as a precise definition (or more exactly a set of definitions) we would find that on one side we have quantitative sentences and on the other we have only qualitative sentences. Indeed the only qualitative relation we have used is equality of degrees of belief; we haven't even used inequalities. When we extend the analysis to real-valued or, on one account, vague degrees of belief, we will need this extra resource.

As I noted above, the idea of defining degrees of belief in this way originates with Ramsey (1929). The first writer to use something like the Equivalence Analysis in a formal theory was Koopman (1940). He argued that probability was primarily a comparative notion; if there are exclusive and exhaustive propositions $p_1, \dots, p_x, \dots, p_y$ each equally likely such that A is as likely as $p_1 \vee \dots \vee p_x$ then A 's probability is x / y , otherwise A doesn't have a numerical probability. To use Koopman's term, in the latter case A is not appraisable. In sections 3.5 to 3.9 we will look at ways of introducing non-appraisable propositions into the theory. Good (1950) uses Koopman's ideas as motivation for the idea that degrees of belief ought obey the probability calculus¹. The theory here uses ideas from all of these writers. The components which are, to my knowledge, original are the use of what I will call models in the third section, and the use of material equivalences at the core of the analysis. It is the last fact which prompted the name *Equivalence Analysis*.

¹ A similar approach is used in Savage (1954). However, he uses 'almost uniform partitions' which are such that any disjunction of $x + 1$ elements is more probable than any disjunction of x elements for any $x \in \{1, \dots, y - 1\}$. The motivation for this is that it makes it more plausible that the p_i can be propositions about real events, rather than dummy propositions as in the theory presented here. Although he starts, like Koopman, with comparative probability, he assumes that for any A, B either $Pr(A) \geq Pr(B)$ or $Pr(B) \geq Pr(A)$. This entails that all propositions are appraisable.

If it helps we can visualise matters by thinking of the various p_i as being the drawing of the i 'th ball from an urn containing y balls, but I'm not sure this helps. In urn cases there is often a temptation to think that the p_i are equally likely because we are ignorant of the way the balls are distributed in the urn. This may be approximately correct in some practical cases, but it seems wrong in general. I am doubtful that we must, or even may, derive precise numerical degrees of belief through ignorance. I certainly don't want to have an appeal to the Principle of Indifference at the core of my definition of degrees of belief. Rather the kind of case I am thinking of is one where our evidence is sufficient to have an equal degree of belief in each p_i . It doesn't take much of a sceptical attitude to deny that physical evidence can ever provide us with this in real world examples. Nevertheless, the situation seems a useful fiction, particularly because in practice our degree of belief in certain types of events (e.g. lotteries) do *approximate* this ideal of equal degrees of belief in all outcomes.

3.2 Outline of Chapter

In section 3.3 I will introduce formal models of probabilistic belief. These are important because they allow us to regain the conclusions of Dutch Book arguments. That is, on the assumption that all an agent's degrees of belief are rational numbers, then it is a coherence requirement that their degrees of belief obey the axioms of the probability calculus. Models are simply sets of propositions closed under entailment, however they are defined on a different possibility space. The core idea behind the models is that if in reality the agent believes A to degree x/y , some proposition of the form $A \equiv p_i \vee \dots \vee p_j$ is true in the model, where there are x disjuncts on the right-hand side. The axioms of the probability calculus fall out as consistency requirements on the model. In section 3.4 we extend this to updating methods for probabilistic beliefs, showing that the Bayesian requirements of Conditionalisation and Reflection can be justified by these models. In that section I also respond to some criticisms of these principles.

There are two complications that can be made. I have assumed in the initial sections that degrees of belief are rational and precise, but in general neither of these restrictions is permissible. In sections 3.5 to 3.9 I look at various ways of dropping the restriction that degrees of belief are precise. The most common way in the literature to do this is to have a person's degree of belief represented by sets of probability functions rather than a single function. Section 3.5 looks briefly at the motivations for imprecision and outlines this approach to representing it. I will then consider two alternatives to this approach.

In section 3.6 I look at a simple alteration of our conditions on models to permit imprecision. In the standard model, when the agent's degree of belief in A is x/y , there is a disjunction D of x elements such that $A \equiv D$ is in the model. In imprecise models, when the agent's degree of belief in A is vague over the interval between x_1/y and x_2/y there are disjunctions D_1 and D_2 of x_1 and x_2 elements respectively such that $D_1 \supset A \supset D_2$ is in the model. It is shown that the resultant theory is equivalent, in a certain sense, to Shafer's theory of belief functions. More importantly the resultant axioms on coherent degrees of belief are *more* restrictive than the approach of section 3.5.

Above I noted that we could think of the p_i as representing the drawing of a given ball from an urn with y balls in it. One way to loosen some of the restrictions on degrees of belief incurred in section 3.6 is to allow there to be multiple urns, not necessarily independent. This approach is analysed in section 3.7. This move doesn't change our earlier conclusions for precise degrees of belief, but it does allow us to drop some of the unwanted restrictions. Although the resultant theory satisfies some proposed axiomatisations of vague probability theory, I will argue that it loses too much structure, and that some of the theorems it fails to prove are ones we ought want. So the only two live options are the family of probability functions approach, outlined in section 3.5, and the Shafer belief functions approach, outlined in section 3.6.

Shafer proposes that his belief functions should not be updated by conditionalisation. This is, I suggest, a mistake, and I'll look at some examples designed to reinforce that belief. Some writers have suggested that this mistake can be remedied. We noted that Shafer belief functions are equivalent to families of probability functions with two particular properties. Although one of these properties is preserved under conditionalisation, the other is not, so it seems there is no coherent way to update Shafer functions.

In section 3.9 I consider Walley's argument that analyses of vague degrees of belief in terms of families of probability functions is bound to give the wrong answer to certain conditionalisation problems. He argues that adopting a certain principle, *conglomerability*, leads to thinking that vague previsions are basic, not probability functions. When applied to general cases the principle of conglomerability is inconsistent, which somewhat vitiates its force when applied to the special case Walley considers.

Finally in section 3.10 I look at extending the analysis of this chapter to real-valued degrees of belief. I say that $Bel(A) = r$ iff $Bel(A)$ is greater than (less than, equal to) y/z for integer y and z whenever y/z is less than (greater than, equal to) r . This way talk of real-valued degrees of belief is eliminated in favour of comparisons between degrees of belief and rationals. An appendix contains some proofs which are left out of the body of the text to ease exposition.

3.3 Precise Models

I'll again adopt the notation $Bel(A)$ for the agent's degree of belief, or credence, in A . The aim of this section is to show that, at least for the special case when Bel takes only rational values, the agent is unreasonable if Bel is not a probability function. As noted in chapter 1, by definition Pr is a finitely-additive probability function iff it is a function from propositions to numbers satisfying:

$$(Pr1) \quad Pr(A) \geq 0;$$

(Pr2) $Pr(T) = 1$;

(Pr3) If $\vdash \neg(A \& B)$ then $Pr(A \vee B) = Pr(A) + Pr(B)$.

If an agent believes, in the traditional sense, A , then they have the same credence in A as they would have in p_1 were they to believe exactly one member of the set $\{p_1\}$ is true. Hence their credence in A will be $1/1$, that is, 1. If they believe $\neg A$ then there will be no value y such that they have the same credence in A as in p_1 and exactly one member of $\{p_1, \dots, p_y\}$ is true. Hence for all y , $Bel(A) < 1/y$. So $Bel(A) = 0$.

Since we are discussing what agents should believe, or what reason requires them to believe, we need to define what concept of reasonableness we have in mind. For these purposes I will simply adopt the standard used in Dutch Book arguments. An agent is reasonable iff their beliefs are closed under entailment and not trivial. This is simply a coherence constraint on reasonableness. For consistency with the literature, I'll use reasonable in this chapter simply to mean probabilistically consistent, or what we might call coherent². This is not meant to imply that all coherent belief sets are reasonable in some strong sense. Nor is it meant to imply that incoherence is unreasonable in an everyday sense.

Above I said that $Bel(A) = x/y$ meant simply $Bel(A) = Bel(p_1 \vee \dots \vee p_x)$ where $\{p_1, \dots, p_x\}$ is a set of propositions such that the agent knows one of them is true and for any i, j $Bel(p_i) = Bel(p_j)$. Put this another way, I could say that $Bel(A) = x/y$ means that the agent has the same credence in A as they would have were they to believe (i.e. believe fully) $A \equiv p_i \vee \dots \vee p_j$, where there are x disjuncts and $\{p_1, \dots, p_y\}$ is defined the same way.

It is clearly no constraint on rationality that if $Bel(A) = x/y$ we can find y equiprobable alternatives such that we believe A is materially equivalent to a disjunction of x of these. However,

² We'll come across a different use of *coherent* later in connection with vague belief functions.

it does seem to be a constraint that we should be able to consistently believe something of this sort. If it were inconsistent to believe that there was any such set as $\{p_1, \dots, p_y\}$ and A is materially equivalent to a disjunction of x elements of this set, something seems amiss.

Indeed, an even stronger constraint than this seems in order. Assume there exists a y such that for all propositions A in a finite field of propositions Γ , $y \cdot Bel(B)$ is an integer. Provided our degree of belief in any proposition in Γ is rational this will be possible for some y . Again assume $\{p_1, \dots, p_y\}$ is a set of propositions such that we know exactly one is true and all are equiprobable, and that no proposition about any p_i is in Γ . Then it seems to be a rationality constraint that it should be possible for any B in Γ , where $Bel(B) = x / y$ to find a disjunction of x elements of $\{p_1, \dots, p_y\}$ such that it is consistent to believe B is materially equivalent to that disjunction. That is, it should be possible to model our probabilistic beliefs about propositions in Γ on $\{p_1, \dots, p_y\}$. Since what it means to say $Bel(B) = x / y$ is to say B is believed to the same degree as such a disjunction, it would be odd if it were inconsistent to say that B is materially equivalent to any of them.

The above is fairly informal, particularly the requirement that no proposition ‘about’ the p_i be in Γ . The following is a more formalised statement of it, followed by proofs that these restrictions are sufficient to show why the degrees of belief should follow the probability calculus.

Assume an agent has a certain set of beliefs, say \mathbf{K} , and certain degrees of belief $Bel(\bullet)$. I want to test for coherence her credences about a certain set Γ of propositions, assuming that all these are rational numbers. A simple coherence constraint, which was argued for above³, is that if $A \in \mathbf{K}$ then $Bel(A) = 1$. Let y be the lowest common denominator of these degrees. Since we are considering her propositional beliefs we can represent each belief as a set of possible worlds, or more generally as a subset of the possibility space. Let W be the set of all worlds in the original

³ And see also section 4.4 for consideration of some recent objections to this proposal.

possibility space. Since \mathbf{K} is closed, it determines, and is determined by, some subset Δ of W . For any proposition A , $A \in \mathbf{K}$ iff Δ is a subset of A . The non-triviality constraint on reasonable belief sets is simply that Δ is non-empty.

Let P be the set $\{p_1, \dots, p_y\}$. I won't presume these are propositions in the sense of being subsets of W . Rather, to allow the agent to consider P , I need to extend the possibility space from W to $W \times P$. Elements of W are possible worlds, e.g. w . Elements of $W \times P$ are ordered pairs, the first element of which is a possible world, and the second element a member of P . Because the p_i are not meant to be about the propositions in Γ , it can be assumed that all points in $W \times P$ are real possibilities. The proposition A on W will be the proposition $\{\langle w, q \rangle: w \in A \ \& \ q \in P\}$ on $W \times P$. We'll write the latter proposition as A^* , and in general use A, B, D for propositions on W , A^*, B^*, D^* for the equivalent propositions on $W \times P$. The proposition p_i in P will be the proposition $\{\langle w, q \rangle: w \in W \ \& \ q = p_i\}$ on $W \times P$. I'll write that as p_i^* . I define conjunction, disjunction and negation of propositions in $W \times P$ in the usual way as intersection, union and complementation.

The aim is then to see if the agent's probabilistic beliefs about Γ can be modelled as propositional beliefs about subsets of $W \times P$. In the model the set of propositions on $W \times P$ which the agent believes is \mathbf{K}^* , which again to satisfy the reasonableness constraint should be closed and non-empty. Hence \mathbf{K}^* determines, and is determined by, a subset Δ^* of $W \times P$. The model should satisfy the following constraints.

$$(1) \quad (S \subseteq P \ \& \ S^* \in \mathbf{K}^*) \rightarrow S = P$$

$$(2) \quad \text{If } Bel(A) = x/y \text{ then } \exists S: ((S \subseteq P \ \& \ |S| = x) \ \& \ (S^* \equiv A^* \in \mathbf{K}^*))$$

In (2) and in what follows I will assume sets of propositions are truth-valued, and I'll simply stipulate that a set of propositions is true iff some element of it is true, false otherwise. This greatly eases the exposition in what follows. Again a proposition unstarred is either on W or P , and starred is the equivalent proposition on $W \times P$. (2) says that \mathbf{K}^* is a model for Bel in the sense that

if A is believed to degree x / y then it is equivalent in \mathbf{K}^* to a disjunction of x propositions such that the agent knows exactly one of y of these is true. (The expression $|S|$ in (2) refers to the cardinality of S). (1) says that the agent doesn't know that some subset of P is true. If this were to occur it would of course be inconsistent with the assumption that each of the p_i is equiprobable. That assumption isn't used anywhere else in the proofs.

We now prove the following lemmas:

- (L1) $\underline{\vee}(p_1^*, \dots, p_y^*) \in \mathbf{K}^*$
(L2) $(S \subseteq P \ \& \ \neg S^* \in \mathbf{K}^*) \rightarrow S = \emptyset$

The notation $\underline{\vee}$ in (L1) refers to exclusive disjunction. There is a difficulty with representing n -place exclusive disjunction, since if we just took $A \underline{\vee} B =_{\text{df}} (A \vee B) \ \& \ \neg(A \ \& \ B)$, we would have the result that $(p_1 \underline{\vee} p_2) \underline{\vee} p_3$ would be true iff one or three of the p_i were true. So I allow exclusive disjunction to be an n -ary connective, written $\underline{\vee}(p_1, \dots, p_n)$, which is true provided exactly one of the p_i is true⁴.

An exclusive disjunction is true iff the disjunction of every element is true and the conjunction of any distinct pair of disjuncts is false. For any two propositions p_i^*, p_j^* , it follows from the definition of $W \times P$ that $p_i^* \cap p_j^*$ is empty. Since Δ^* is not empty, it follows that $\Delta^* \not\subseteq (p_i^* \cap p_j^*)$. Hence $p_i^* \ \& \ p_j^* \notin \mathbf{K}^*$. Again from the definition of the p_i^* , it follows that $p_1^* \cup \dots \cup p_y^* = W \times P$. Hence $\Delta^* \subseteq (p_1^* \vee \dots \vee p_y^*)$, so $p_1^* \vee \dots \vee p_y^* \in \mathbf{K}^*$. This proves (L1).

(L3) follows directly from (L2), since if $S \subseteq P$, then $\neg S$ will be P / S , which is also a subset of P . Since $\neg S^* \in \mathbf{K}^*$, by (L2) $\neg S = P$, so S must be empty.

⁴ McCauley (1993) defines all of the connectives as n -ary, partially to have a greater parallel between formal and natural language and partially to resolve this problem with exclusive disjunction.

These rules are sufficient to prove that the following are constraints on reasonable degrees of belief for any propositions A, B in Γ . (The proofs are in the appendix to this chapter)

Theorem 3.3.1 If Bel can be modelled by \mathbf{K}^* satisfying (1) and (2) above and is defined for all propositions in Γ , then for all A, B in Γ :

- (T1) If $\vdash A$ then $Bel(A) = 1$
- (T2) If $\vdash \neg A$ then $Bel(A) = 0$
- (T3) $Bel(A) + Bel(B) = Bel(A \vee B) + Bel(A \& B)$
- (T4) If $A \vdash B$ then $Bel(A) \leq Bel(B)$
- (T5) $0 \leq Bel(A) \leq 1$
- (T6) If $Bel(A) = x / y$ and $Bel(A) = z / y$ then $x = z$

If we can prove (T1) to (T6) then it follows that Bel is reasonable *vis a vis* Γ iff Bel obeys the axioms of the probability calculus with respect to those propositions. Generally Bel is reasonable (or at least coherent) if it is reasonable *vis a vis* any finite set Γ of propositions.

Now it simply falls to us to show that these requirements ensure that Bel is a probability relation. By (T6) Bel is a function from propositions to numbers (i.e. it is uniquely valued). By (T5) it satisfies (Pr1), by (T1) (and (T6)) it satisfies (Pr2) and by (T2) and (T3) it satisfies (Pr3). Hence it is a probability function.

3.4 Updating Precise Models

Dutch Book arguments have been presented to show that we should obey certain dynamic principles, such as Reflection and Conditionalisation. These are commonly referred to as Dutch Strategy arguments, and these seem most at risk from my arguments in section 2.3. We should look at whether we can reach the same conclusions as they do using the Equivalence analysis of degrees of belief. To do this I'll first look at what similar results to Reflection and Conditionalisation can be found concerning absolute beliefs.

Again the only norms I am adopting here are norms of dynamic coherence. I take it that if there is no possible world in which an agent can believe truly a set of propositions \mathbf{K} then \mathbf{K} is not a statically coherent belief set. However, since we are studying epistemic dynamics we have to extend this rule. So I adopt as a coherence constraint that holding a belief such that my holding that belief entails I have *or will have* a false belief shows that I am unreasonable. There are also closure constraints on rationality. Assume that Δ is the set of all worlds in which we have all the current beliefs we actually have and all our current and future beliefs are true. The first coherence requirement is that Δ be non-empty. As above, I will also insist that reasonableness requires that for every proposition A true in every world in Δ , we believe A .

Assume I believe that tomorrow I'll have a false belief. ('Tomorrow' refers simply to an arbitrary future time). If this is true I will, perforce, have a false belief. Alternatively, this belief will be false, so again I will have a false belief, this one. Hence I am guaranteed to have a false belief, so by the coherence constraint I am incoherent, or unreasonable. Hence it is unreasonable to believe that tomorrow I'll have a false belief. Similar constraints apply to the belief that I now believe something false. It might be thought that this implies closure is not only unnecessary for reasonableness, it is incompatible with it. However, we get more plausible results when we restrict our attention to particular beliefs.

Assume I believe A and believe that tomorrow I'll believe $\neg A$. Since it can't be that A and $\neg A$, it can't be that my beliefs on each day are true, so one of them must be false. So I must either have a false belief, and hence be unreasonable, or believe I'll have one, which is unreasonable for the reasons we noted in the previous paragraph. Hence again I am unreasonable. Alternatively, assume simply that I believe that tomorrow I'll believe A . Then in all worlds in Δ tomorrow I truly believe A , hence all worlds in Δ are A -worlds. So by closure I should believe A .

The argument for Reflection is just the probabilistic version of that argument. Reflection is the principle that if I believe I will have credence r in A I must now have credence r in A . Assume that today my credence in A is x/y , and I believe that tomorrow it will be z/y , where $z \neq x$. If tomorrow my credence isn't z/y , I currently have a false belief, so tomorrow my credence must be z/y . (Note that by 'it must be that B ' I just mean all worlds in Δ are B -worlds). Now our coherence requirements on credences mean that for any proposition B which we believe to degree v/y , we can consistently add to our belief set that $B \equiv p_1 \vee \dots \vee p_v$. (Where there is little possibility of confusion I drop the $*$ notation.) If we are assessing our beliefs about more than one proposition, the only plausible requirement is that we could add $B \equiv S$ for some $S \subseteq P$ with $|S| = v$. But when just one proposition is being considered this stronger requirement looks plausible enough.

Now, our assumptions lead us to believe that today I believe $A \equiv p_1 \vee \dots \vee p_x$, and tomorrow I will believe $A \equiv p_1 \vee \dots \vee p_z$. Since all my beliefs will be true, this means it must be that $p_1 \vee \dots \vee p_x \equiv p_1 \vee \dots \vee p_z$. If $x > z$, this means it must be that $\neg(p_{z+1} \vee \dots \vee p_x)$. Hence by the closure requirement I should believe $\neg(p_{z+1} \vee \dots \vee p_x)$. But (L3) above still applies, showing that I should only believe the negation of a disjunction of elements of P if that disjunction had zero-elements. So it can't be that $x > z$. Similarly if $z > x$ we would end up being committed to $\neg(p_{x+1} \vee \dots \vee p_z)$, which again is unreasonable by (L3). So reason requires that $x = z$. That is, if we have any beliefs about our future credence in A , we should believe we will believe it to just the degree we currently do.

Two comments on this proof. First, I used the strong requirement that if our degree of belief in A is x/y we should be able to consistently add $A \equiv p_1 \vee \dots \vee p_x$, rather than the weaker requirement that we be able to add $A \equiv S$ for some suitable S . If we adopt just that requirement the proof doesn't go through. For it doesn't follow that if adding $A \equiv S$ implies that tomorrow we will believe of some p_i or other that it is false that we are irrational. To see this, let A be the proposition that the

coin I am about to toss will land heads. Now my credence in A is $1/2$. That is, I could add $A \equiv p_1$, with $y = 2$. However, I believe that tomorrow I'll either believe A or disbelieve A . So if I had added $A \equiv p_1$ I believe that tomorrow I'd either believe $\neg p_1$ or $\neg p_2$. Since this isn't irrational, that I will come to disbelieve some p_i or other can't be irrational, and we can't plausibly strengthen our coherence constraints to make it so.

The other comment is that people have often misinterpreted coherence constraints. It might be perfectly reasonable, given that I am confident I won't be completely coherent tomorrow, to believe I will falsely believe A . The Reflection criteria are necessary but insufficient criteria for full rationality. There is no reason to suspect, and certainly nothing in the above proofs, to justify the idea that Reflective agents are *ceteris paribus* more reasonable than those who are not. Compare the following example. Only students who don't answer an even number of questions on the exam will get a perfect result. It doesn't follow from this that students who don't answer an even number of questions are *ceteris paribus* better students than those who make an odd number of mistakes. Nor does it follow that a student, noting she has failed to answer an odd number of questions but unable to make any more progress, should take steps to ensure she leaves out an even number of questions, particularly if this involves leaving out a good answer. This doesn't invalidate the original claim that not answering an even number of questions is a necessary condition of perfection. Nor does the fact that taking steps to be Reflective would be positively irrational in some circumstances invalidate the claim that Reflection is a criterion of full rationality.

The other constraint which is usually thought to be proven by Dutch Strategy arguments is Conditionalisation. As van Fraassen has pointed out, in its standard form this argument is noticeably weaker than other Dutch Book arguments; indeed considerably weaker than van

Fraassen's own argument for Reflection (van Fraassen 1989: 173 - 6)⁵. Assume that my current degree of belief in A given B is 0.2, but if I were to discover B (and presumably nothing else) my credence in A would be 0.3. However, I do not know this, I believe falsely that I would be a good Bayesian and conditionalise. Then someone who knows all this can offer me bets that will lead to me losing money. But this doesn't prove any incoherence. After all, assume my credence in A is 0.2, but A is false. Then someone can offer me a bet on A for 10 cents, which I'll buy and lose. The fact that I'm prepared to take losing bets doesn't show irrationality if it requires more information than I have to show that the bets are losing ones. So the Dutch Book argument to show that not being a conditionaliser is irrational relies on the presumption that I know, or at least ought know, what my degree of belief in a given proposition would be under any circumstances. I clearly don't have this information, and I don't see any particular reason why I should be required to do so. Hence the traditional Dutch Book argument fails, because the Dutch Bookie can only make a book if she has more knowledge than I either have or am epistemically obligated to have, and I'm only irrational if she can make a book without more knowledge than I have or ought have.

van Fraassen showed that we could revive the Dutch Book argument against a person intending to follow a strategy other than Conditionalisation. Like all dynamic Dutch Book arguments, his appears to be faulty for the general reasons I have gone through; however it is at least acceptable by its own standards. He concluded that it could be rational to not conditionalise, as long as one did it capriciously rather than in accord with a pre-ordained strategy. This conclusion has been questioned by some who thought it paradoxical that a course of action could be irrational if carried out deliberately, but rational if carried out capriciously. There is certainly a ring of paradox about it!⁶

⁵ An argument against requiring conditionalisation based around this fact is also found in Howson (1993).

⁶ See, for example, Green and Hitchcock (1994: 321).

My arguments for Conditionalisation do not rely on Dutch Book considerations. Rather, they rely on the Equivalence Analysis of degrees of belief and the models used to show that degrees of belief ought obey the axioms of the probability calculus. Before we start on them however, we need to consider briefly the idea of modifying standard belief sets. There is a large literature on this, most of which is concerned with the difficult problem of how to retract beliefs. I only want to consider a simple problem, how to add a propositional belief to a coherent set of propositional beliefs.

As we noted above, if an agent has propositional beliefs \mathbf{K} and \mathbf{K} is closed under entailment, then \mathbf{K} will determine and be determined by a set Δ of possible worlds. There seems to be something of a consensus that the way to add a proposition, say B , to a set of beliefs \mathbf{K} , is for the new belief set to be the closure of $\mathbf{K} \cup \{B\}$. Or in other words, for the new belief set to be that set determined by the set of worlds $\Delta \cap B$. I will adopt this conclusion in what follows. There could be a slight complication if \mathbf{K} contains $\neg B$. Since this would involve difficult questions about revision of belief sets, an issue on which there is no agreement amongst theorists, I'll assume $\neg B$ is not in \mathbf{K} and more generally that $Bel(B) > 0$.

Assume $Bel(A \& B) = x / y$ and $Bel(B) = z / y$, where $z > 0$. I have to prove that when we add B to \mathbf{K} the only rational new value of $Bel(A)$ is x / z . If Bel and \mathbf{K} are rational then there is some \mathbf{K}^* as outlined above containing $S^* \equiv B^*$ for some S such that $|S| = z$. Although there is no requirement that S be $\{p_1, \dots, p_z\}$ we can rename the p_i such that this is the case. After all, there is no distinction between the various elements of P . For convenience, I'll call the set $\{p_1, \dots, p_z\}$ S_z . There will also be a set S such that $(A \& B)^* \equiv S^*$ is in \mathbf{K}^* , $|S| = x$ and $S \subseteq \{p_1, \dots, p_z\}$.

Now \mathbf{K}^* is a propositional belief set. So we can see what beliefs that agent should have if they started with that set and added B , or equivalently, added B^* . The new set, which we'll name \mathbf{K}_B^* , is simply the closure of $\mathbf{K}^* \cup \{B\}$. Since \mathbf{K}^* contains $B^* \equiv S_z^*$ and B , it contains S_z^* . Let the new

set of beliefs and degrees of belief the agent has after coming to believe B be called \mathbf{K}_B and Bel_B respectively. Since \mathbf{K}_B^* contains $\bigvee(p_1, \dots, p_z)$, if for any proposition D , there is a set S such that $S \subseteq S_z$ and \mathbf{K}_B^* contains $D \equiv S$, then $Bel_B(D) = |S| / z$. If the agent is reasonable then by (T6) any way of calculating $Bel_B(D)$ will give the same answer.

Now we know that there is some S such that $|S| = x$ and \mathbf{K}^* contains $S^* \equiv (A \& B)^*$. Hence \mathbf{K}_B^* also contains $S^* \equiv A^*$, since it is a superset of \mathbf{K}^* . And we also know $S \subseteq S_z$, as $(A \& B) \vdash B$. So it follows that $Bel_B(A \& B) = x / z$. Since \mathbf{K}_B contains B , we know that $Bel_B(B) = 1$, and by (T4) $Bel_B(A \vee B) = 1$. Hence by (T3) $Bel_B(A) = Bel_B(A \& B) = x / z$ as required.

In summary, the argument is that given \mathbf{K} and Bel represent the agent's original belief state, the agent's beliefs must be as they would be as if the agent believed \mathbf{K}^* . We know how to reasonably amend \mathbf{K}^* by adding B , since how to add beliefs to a non-probabilistic belief state is non-controversial. Hence whatever the new agent's beliefs are, say \mathbf{K}_B and Bel_B , we know that they must be capable of being modelled by \mathbf{K}_B^* , which we can determine. From this and the simple assumption that $Bel_B(B) = 1$ it follows that $Bel_B(A) = Bel(A \& B) / Bel(B)$.

The conclusion of this argument is *stronger* than the conclusion of the parallel Dutch Book argument. As I noted above, that argument could only prove that it was irrational to adopt any strategy other than Conditionalisation. However, this argument shows that it is irrational not to adopt the strategy of Conditionalisation.

3.5 Introducing Imprecision

So far I have assumed that all degrees of belief are precise, or in Koopman's term, appraisable. It is time to relax that assumption. I will look in detail at the arguments for this move in chapter 5, but I will introduce them here to motivate what follows. Like precise degrees of belief, imprecise degrees of belief should be capable of being modelled. This I take as a coherence constraint. I will introduce my preferred method of modelling and then look at two advantages of it. The

section concludes with a look at the history of this approach to imprecision, and some possible reasons for moving away from it. This leads naturally to the alternatives outlined in subsequent sections.

3.5.1 *Why Be Imprecise?*

There are three arguments for the conclusion that we shouldn't require that reasonable degrees of belief be precise. The first is from introspection, the second from the possibility of ignorance and the third from the possibility of rational disagreement.

When we look at our credences in various propositions, it seems highly plausible that these are not precise. My credence in *Oswald killed JFK* is not a precise number; that is, there is no fair lottery with y tickets such that my degree of belief in this proposition equals my degree of belief that one of the first x tickets will win.⁷ I might be wrong, but I don't think this is because of a failure of rationality on my part. Rather, I have too little clear information on the Kennedy assassination to form a precise credence.

If we believed a betting analysis of credences we could follow a process derived by Borel to measure, to any degree of accuracy we wanted, my degree of belief in this proposition. For arbitrarily large y we could keep offering me a choice of bets on Oswald being the killer or on one of the first x tickets winning. As long as I choose the Oswald-bet, we increase x until I first choose that bet. In response to Borel I can firstly run through my earlier objections to the betting analysis, but I can make one extra point. After a while (say after x grows beyond $0.4y$ or thereabouts) my choices would simply be arbitrary, and would not directly reflect my credences.

⁷ If we are allowing real degrees of belief we have to say this more carefully in terms of sequences of fair lotteries. See section 3.10 for the technical details.

The assumption that my credences determine what I would do in all circumstances is just the completeness assumption that is at issue here.

If we allow credences to be imprecise we have a good way to represent complete ignorance. Classically, ignorance was represented using Laplace's Principle of Indifference. If we didn't know anything about what would happen, we allocated equal probability to each possibility. Unfortunately when applied indiscriminately this led to inconsistency. Even when applied consistently it led to results which really were absurd.

For example, let q be the proposition that there is intelligent life somewhere else in our galaxy. Put some numbering on all the stars in the galaxy (there's only finitely many of them), and for each star i , let q_i be the proposition that there is intelligent life on some body orbiting that star⁸. We are presumably completely ignorant about q , and about each q_i , so by the Laplacean principle we assign credence $1/2$ to each proposition. To see this more clearly, note that when we are considering whether or not it is the case that q , there are two possibilities: $q, \neg q$. The Laplacean principle says to assign equal credence to each possibility, which in this case means assigning $1/2$ to q and $1/2$ to $\neg q$. However, the same reasoning applies to each q_i . Since each q_i entails q , it follows by simple applications of the probability calculus that we believe absolutely that if there's intelligent life orbiting some other star in the galaxy there is intelligent life orbiting all of them. This is hardly the kind of thing we ought be able to infer from ignorance. When we allow credences to be imprecise on the other hand, we can say that for q and each q_i that our credence in it is vague over the interval $[0, 1]$. We will have to say something about what this means, but it turns out not to have any absurd consequences.

⁸ I intend 'orbiting' to be read widely enough that the moon does orbit the sun, which indeed it does in a way.

These two arguments are fairly well known, however the third argument is new. In large part this is because it only arises when we analyse probability, as it has been analysed here, as reasonable credence. It is a commonly observed fact that people with the same evidence, or the same relevant evidence, have different credences in a proposition. This can happen even when it seems plausible to say that each is acting reasonably. It would be a benefit for a theory of probability if it licensed this conclusion. That is, if it allowed us to say that reasonable people with the same evidence could, at least some of the time, have divergent credences in a proposition. And it turns out if we allow credences to be imprecise we can say exactly this.

These three arguments are enough, I hope, to motivate an investigation into the technical properties of imprecise credences. I leave the detailed arguments until later, because I want to use some of the technical apparatus developed here in presenting these arguments.

3.5.2 *The Many Models Approach*

When all the credences of a coherent agent are precise, they can be represented by a single probability function Pr . On my preferred theory, when they are imprecise they can be represented by a family of probability functions P . What it means to say that her credence in q is vague over an interval $[\alpha, \beta]$ is that this is the smallest interval such that for all $Pr \in P$, $Pr(q) \in [\alpha, \beta]$. Following van Fraassen (1990) I will call P her ‘representor’.

If an agent’s credences can be represented in this way there is a trivial sense in which they can be modelled. Every Pr in P can be modelled, as shown in section 3.3. Hence we can simply say that the model for each Pr is a possible model for the agent. One way of putting this is to say that the correct model for the agent is vague over these possible models. A better way is to say that there are many models for this agent. So I call this the Many models approach.

Apart from arguments from authority there are three reasons for taking this line. First, by using supervaluations it allows us in a sense to have the best of both worlds. In standard cases supervaluationist approaches allow us to keep classical logic, while still saying that the truth value of some propositions is vague. Here it allows us to keep the probability calculus in all its power, while allowing us to say that the probability of some propositions is vague. This will be developed further in chapter 5.

The second reason follows from some technical work done by Smith (1961) and Williams (1976)⁹. A *gamble* is a bet that has a certain (finite) payout in each possible world. The bets we have been looking at so far pay $\$1 - \delta$ at some worlds and $-\delta$ at others, where δ is the price of the bet. However, these more general gambles can have all sorts of different payouts. This allows us to talk, for example, about the addition of gambles. We can think of gambles as functions from worlds to reals. Then in the usual way we add functions, we can add gambles. So for all worlds w , and gambles $x, y, x + y(w) = x(w) + y(w)$. Now let D be the class of strictly desirable gambles¹⁰. We adopt as coherence axioms that D is closed under addition and positive scalar multiplication. We also assume that D includes all gambles which always take positive values and no gamble which always takes negative values.

Now we can, in the usual way, work out the expectation value of a gamble according to a probability function Pr . So we can work out the set of expectation values that a gamble takes

⁹ The exposition here follows Walley (1991).

¹⁰ I include the term ‘strictly’ because if we think that we will only accept desirable gambles then vagueness leads to a kind of dynamic incoherence, set out in chapter 10. When we add ‘strictly’ it should be clearer that we don’t think that D exhausts the class of gambles we would accept if offered. Put another way, vague decision theory is only dynamically coherent if there are gambles we are neither disposed to accept nor reject; calling the class of gambles we are disposed to accept the ‘desirable’ gambles may lead to the misapprehension we’re disposed to reject the rest.

according to every element of a family \mathcal{P} . It turns out that whenever D obeys the above properties, there is a family \mathcal{P} of probability functions such that a gamble x is strictly desirable iff its expectation value is positive according to every element of \mathcal{P} .

This result can be put the following way. If the agent's credence in every proposition in a field are not such that they can be represented by a family of probability functions, then (assuming a simple decision theory) that agent's decision-making will be incoherent in this sense. There is a set of gambles $\{y_1, \dots, y_n\}$ such that each y_i is strictly desirable, but $y_1 + \dots + y_n$ is not strictly desirable. This looks to be a bad consequence. Note that this is a much stronger requirement than the usual Dutch Book requirement on credences. An agent who is incoherent in this sense need not be susceptible to a Dutch Book. However, whether the individual considers some gambles strictly desirable will be dependent on the way the gambles are presented, not the contents of those gambles.

For example, assume that propositions q_1 and q_2 are disjoint. Assume that for q_1 , q_2 and $q_1 \vee q_2$ the agent's credence is vague over $[0.25, 1]$. Let y_i be a gamble which pays 80 cents if q_i , -20 cents otherwise, for $i \in \{1, 2\}$. Then the agent will find y_1 and y_2 strictly desirable, but not $y_1 + y_2$. There is no way this agent can have a Dutch Book made against them, there is no set of strictly desirable bets which have in sum a uniformly negative payout, but something about their degrees of belief seems defective. It is this 'something' which Smith and Williams are interested in, and it turns out that were their degrees of belief represented by some family of probability functions this incoherence would go away.

As we said in section 2.5, the betting analysis might not be a successful analysis for degrees of belief. However, it is a useful analogy, one of the best analogies we have. And as we've seen already in this chapter, most of the controversial conclusions argued which are usually argued for

by Dutch Books arguments are correct. So I take Smith and Williams's arguments to be persuasive but not necessarily compelling arguments for the Many models approach.

3.5.3 *Disjunction And Inequalities*

There is a third argument for taking this approach to representing or modelling vague degrees of belief. Let q be some proposition in which my degree of belief is vague over an interval of width greater than $1/y$ for some finite integer y , and smaller than $(1/y, (y - 1)/y)$. These are fairly minimal constraints. Let q_1 be the proposition that ticket #1 in a fair lottery with y tickets will win. Assume that the fairness of the lottery entails that it is independent of q which ticket will win.

This seems enough for it to be intuitively clear that my $Bel(q \vee q_1)$ is greater than $Bel(q)$. This is despite the fact that the intervals over which the two are vague overlap. We now have three options. First, we could bite the bullet and deny that $Bel(q \vee q_1)$ really is greater than $Bel(q)$, introspective evidence notwithstanding. Since I take introspective qualitative data as my fundamental given, this seems implausible.

Whenever $Bel(A)$ is vague over $[\alpha, \beta]$, let $Bel_*(A) = \alpha$ and $Bel^*(A) = \beta$. The second option then is to say that $Bel(A)$ is greater than $Bel(B)$ simply means $Bel_*(A) > Bel_*(B)$, or alternatively it means that $Bel^*(A) > Bel^*(B)$. There are some difficulties with updating under this approach, but as using my preferred updating rule (i.e. conditionalisation) to cause problems would probably be begging the question I won't stress this point¹¹. Rather I want to note some difficulties this approach has with negation. We showed above that in the precise case if $Bel(A) > Bel(\neg A)$ then $Bel(A) > 1/2$. Since the results we obtained there followed so simply from the Equivalence Analysis, I prefer approaches to vague credences which keep as many of these results as possible. However, that approach to inequalities does not preserve this result.

¹¹ See however the discussion in chapter 9 of the analogous decision-rule, called Maxi.

Assume $Bel(A)$ is vague over $[0.45, 0.6]$. I presume this means $Bel(\neg A)$ is vague over $[0.4, 0.55]$. I leave until chapter 8 discussion of approaches which might not have this result, but it might be noted that this holds (suitably interpreted) for almost all the approaches to vagueness on the market.¹² On the approach to inequalities advocated here, it follows that $Bel(A) > Bel(\neg A)$, but it is not the case that $Bel(A) > 1/2$. Without a motivation for accepting this, I take it that this approach to inequalities fails.

The third option, and it seems the only plausible one, is to say that we have to look at more than the bounds of $Bel(A)$ and $Bel(B)$ to determine what inequalities hold between them. And under the Many Models Approach we do just this. We say $Bel(A) > Bel(B)$ for an agent represented by \mathcal{P} iff for all Pr in \mathcal{P} , $Pr(A) > Pr(B)$. A similar definition is given for equality of credences, from which it follows that several credences will be incomparable.

Again, this isn't a totally compelling argument for the advocated approach. There might be other ways at looking at the 'structure' of the intervals other than the family of probability functions approach. And there might be ways of saving the second option, say by motivating the rejection of the theorem on which I relied. However, again the argument developed in this subsection seems to be at least persuasive.

3.5.4 History of This Approach

The idea of representing vague degrees of belief by sets of precise probability functions has gained dramatically in prominence in recent years. The views of three of its modern proponents (Levi, van Fraassen and Jeffrey) are examined in detail in chapter 7. This subsection looks at its prehistory. The motivation stretches back to the distinction between *risk* and *uncertainty* drawn by the American economist Frank Knight in his (1921), and the non-numerical probabilities

¹² This includes approaches where all we specify is a lower bound for $Bel(A)$. Usually the upper bound will be determined in effect by $1 - Bel(\neg A)$.

promoted in Keynes's (1921a). However, neither theorist could be said to have had this particular formal representation of their informal ideas in mind¹⁵.

Usually credit for the sets of probability function approach is given to the statistician Cedric A. B. Smith. His (1961) expressed many ideas which were important for the subsequent development of the theory. In particular the coherence constraints mentioned above are originally his. In discussion on that paper, I. J. Good claims the idea is quite similar to his 'black box' approach, which was first set out in his (1950), particular on page 32 of that book. However, inspection of that book, and in particular that page, reveals nothing which could be construed as a statement of the theory. These matters are always a little vague, but it seems to me the first statement of Good's black box theory is his (1962). Whenever Good first had this idea, which in all likelihood was well before Smith's paper, my reading is that he was beaten into print.

As far as priority goes none of this is, however, particularly relevant. Credit for the idea is, I think, best awarded to the econometrician Gerhard Tintner and economist A. G. Hart. Tintner's (1941) is a brief attempt to capture formally some of the informal ideas that drove Knight's work. It isn't particularly clear but I think it does just enough to be given priority. This summary of his ideas is from his opening paragraph; hopes that its lack of clarity are just caused by its being a concise summary are dashed by inspecting his text.

Subjective risk deals with the case in which there exists a *probability distribution* of anticipations which, however, is itself known with certainty (probability one). Subjective uncertainty assumes that there is an *a priori probability of the probability distributions* themselves, i.e. a

¹⁵ In Chapter 10 I argue that this is a formal representation of Keynes's ideas, though this might be contentious. Knight is a bit harder to classify because he seems to be drawing a different distinction in different chapters of his book. See Schmidt (1996).

distribution of the probability distributions (Tintner 1941: 298, italics in original).

It isn't clear in Tintner why we oughtn't reduce the second-order probability distributions to a first order one. Indeed, if we take his probability distributions to be hypotheses about the objective chance of a propositions, and his *a priori* probability to be degree of (rational) belief in these hypotheses, this is precisely what ought be done. Perhaps were he trying to write a philosophy paper and not an econometrics paper he would have been clearer on these issues. Or perhaps he just made a dumb error. It's hard to say from the text which is correct, and charity demands that we at least not accept the latter.

Hart's (1942) is somewhat clearer on these points. He thinks that there will be occasions under which we should just reduce the set of distributions to a single distribution, in particular where a decision is needed now. However, when we don't need to act immediately on our partial beliefs then having vague degrees of belief is just the reasonable step of keeping one's options open. I will return to this possible connection between vague degrees of belief in chapter 11 and preference for delaying decision.¹⁴ Hart also notices that vague degrees of belief serve as an explication of at least something Knight was trying to capture with the distinction between risk and uncertainty.

'Risk' is taken to denote the holding of anticipations which are not 'single valued' but constitute a probability distribution having known parameters. 'Uncertainty' is taken to denote the holding of anticipations

¹⁴ When I formally set out the idea I associate holding money, or gambles with a more stable money value, with delaying action. Hence I won't need an implausible distinction between acts and omissions.

under which the parameters of the probability distribution are themselves not single valued (Hart 1942: 110).

Neither Tintner's nor Hart's work receives a great deal of attention, and the papers of Smith and Good from 1961 and 1962 respectively appear to have eventually had greater effect. That effect though was somewhat delayed. Indeed to bring the story up to the 1980s the only papers that need to be mentioned are two by Peter Williams. His (1976) was a formal development of some of the ideas in Smith, which from our perspective is notable for being the first to note the suitability of supervaluational semantics to this account of probability. And his (1978) review essay of Shafer (1976) introduced the set of probability functions approach to a much wider audience.

3.5.5 *What might go wrong?*

It might be wondered why I have kept such a tight connection with precise degrees of belief as we stride into the brave new world of vagueness. Why not represent our vague degrees of belief as a single vague model rather than as a set of precise models? The arguments above presented some sort of case for this path, but it's hardly irrefutable.

The conclusion of this chapter will be that the reasons we don't move to a single vague model are that the only plausible way of doing so is (i) *more* restrictive than the approach advocated here and (ii) incapable of being coherently updated. There is no plausible single model which allows an agent to have as its representor any family of probability functions. In the next section we will look at a type of single model which is plausible, but which puts unjustified restrictions on the types of families which are rationally permissible. In section 3.7 we'll look at a way of loosening these restrictions. Unfortunately, in doing so we lose some restrictions we should (I hope) like to keep. For better or worse, the most plausible way to represent vague degrees of belief is by reference to these precise degrees of belief.

3.6 The Single Model Approach

In the previous section we set out one possible paradigm for the representation of vague beliefs, what I call the Many Models Approach. In this section we will look at a different approach, which looks like a more natural generalisation of the approach to representing precise degrees of belief. It turns out this approach is *more* restrictive than the Many Models Approach, so we will look at whether the extra restrictions can be justified. I will conclude, somewhat tentatively, that they are not. However, I don't think this approach is so implausible that we should abandon it altogether, and in section 3.8, which is on updating, we will return to it.

3.6.1 The Model

In section 3.3 we set out a method for modelling an agent's beliefs about a set of propositions Γ such that the agent had a precise degree of belief in each of them. The plausibility of such modelling followed from the definition of degrees of belief set out in section 3.1. The idea is that there is a closed set of propositions \mathbf{K}^* on a possibility space $W \times P$ which satisfies the following two conditions:

- (1) $(S \subseteq P \ \& \ S^* \in \mathbf{K}^*) \rightarrow S = P$
- (2) If $Bel(A) = x/y$ then $\exists S: ((S \subseteq P \ \& \ |S| = x) \ \& \ (S^* \equiv A^* \in \mathbf{K}^*))$

The motivation for (2) was that if the agent had credence x/y in A , their credence in A is just the same as it would be were they to believe $A \equiv S$ for some proposition S in which their credence is, by definition, x/y . However, there is no reason in general for saying that the agent will have a precise credence in every proposition. Hence (2) should be replaced with (3) and (4).

- (3) Iff $Bel(A) \geq x/y$ then $\exists S: ((S \subseteq P \ \& \ |S| = x) \ \& \ (S^* \rightarrow A^* \in \mathbf{K}^*))$
- (4) Iff $Bel(A) \leq x/y$ then $\exists S: ((S \subseteq P \ \& \ |S| = x) \ \& \ (A^* \rightarrow S^* \in \mathbf{K}^*))$

So as not to increase the confusion, I'll here use ' \rightarrow ' for material implication. It might be noticed that (3) and (4) entail (2), so the new axioms are stronger than those for (2). This seems right; even if the agent needn't have precise degrees of belief in every proposition, for every proposition where their credence is precise then (2) ought be satisfied.

I haven't required here, as I did for the precise models approach, that y be related in any interesting way to the end-points of the intervals over which the agent's credences are vague. If (3) and (4) are plausible requirements, it seems they are plausible for any value of y . To make it clear which y we are using, we will use P_y for integer y to represent the set $\{p_1, \dots, p_y\}$. The claim is then for any reasonable epistemic state and any integer y there is a model on $W \times P_y$.

The idea behind (3) is that if $|\mathcal{S}| = x$ then for any proposition B such that $S^* \equiv B^*$ is in the model, $Bel(B) = x / y$. As I showed above, for precise models it is a theorem that whenever $B \rightarrow A$ is believed fully, it is a coherence requirement that $Bel(B) \leq Bel(A)$. I assume (without strictly proving it) that this holds for imprecise credences too, and this is sufficient for the reverse direction of (3). The forward direction follows from the fact that \mathbf{K}^* is a model of the agent's credences. If the largest S satisfying $S^* \rightarrow A^* \in \mathbf{K}^*$ had a cardinality $x - \delta$, then it would seem the agent's credence in A should, according to the model, be $(x - \delta) / y$. Hence either the model is inaccurate, or the agent's beliefs are incoherent in some way. A similar justification applies to (4).

3.6.2 *The Three Prisoners Problem*

To see how this might work, consider the well-known Three Prisoners Problem (TPP). There are three prisoners, a , b and c , scheduled to be executed. The governor has decided to reprieve one of them, chosen at random by a fair chance mechanism. He makes his decision, and notifies the guards of it. However, he doesn't tell the prisoners. Prisoner a , desperate for information, asks a guard if he is to be spared. The guard won't say, so a asks him for the name of one of the other

prisoners who isn't going to be spared, and the guard agrees to answer this. We'll look in section 3.8 about how a should react to the guard's answer, but here we'll just show how a 's beliefs can be modelled to demonstrate the Single Model Approach¹⁵.

Assume, for the sake of simplicity, that a knows the guard to be perfectly reliable. Then there are four possible outcomes. Let the proposition $c_f c_k$ for prisoners c_f, c_k mean that c_f is to be reprieved and the guard says c_k is to be executed. The original possibility space W is $\{ab, ac, bc, cb\}$. If b or c are to be spared the guard has no choice about what to say, so $Bel(bc) = Bel(cb) = 1/3$. (I am assuming here that Lewis's Principal Principle holds: when we know the chance of a proposition q is x , we ought to believe q to degree x .) However, if a is to be reprieved, the guard has a choice about what to say, and a has no way of knowing what he will do. In Gardner's formulation of the problem the prisoner asks the guard to decide in this case what to say by tossing a coin, but my prison guards aren't *that* cooperative. Hence, a 's credence in ab might be vague over $[0, 1/3]$. However, the proposition $ab \vee ac$ is equivalent to the proposition ' a is reprieved', which a knows has chance $1/3$. Hence his credence in $ab \vee ac$ should be $1/3$.

¹⁵ This problem is usually credited to Gardner (1961), though he introduces it as a problem that was 'doing the rounds' at the time (pg 227). If we were being fussy about priority we'd have to say that the problem is a variation on a problem found at the beginning of Bertrand's classic (1889: 2-4). In that problem we have three boxes containing two coins each. The first contains two gold coins, the second a gold and a silver coin and the third contains two silver coins. A box is chosen by a fair chance mechanism and then a coin is drawn from that box by an unknown mechanism. The coin drawn is silver. What is the probability that the other coin in the box is gold? (Bertrand asks the complementary question: What is the probability the other coin is silver?) As in the problems I discuss, the popular but incorrect answer is $1/2$. If we say the coin selection mechanism is fair, then a simple application of Bayes's Rule gives the answer $1/3$. If we leave it unknown I take it the probability of drawing a gold coin is just the same as a 's probability of being reprieved after hearing the guard's news.

We can represent all of a 's beliefs as a model on $W \times P_3$. The model \mathbf{K}^* contains the closure of the following formulae: $p_1 \equiv bc$, $p_2 \equiv cb$, $p_3 \equiv ab \vee ac$. To see how (3) and (4) are satisfied for $A = ab$, note that in the model we have $\perp \rightarrow ab$, $\perp = \emptyset^*$ and $|\emptyset| = 0$, so \emptyset is the S for (3). For (4), we have $ab \rightarrow \{p_3\}^*$ in the model and $|\{p_3\}| = 1$. Hence $\{p_3\}$ is the S in (4). We can show how ac is satisfied in a similar way.

3.6.3 Shafer Functions

To show that every set of vague credences which can be represented under this approach can be represented using the Many models approach, I will show that this approach is equivalent to one advocated by Shafer (1976). His book built on some earlier suggestions of Dempster (1967, 1968), so the functions of the theory are usually referred to as Dempster-Shafer functions, or just Shafer functions. They are also called ‘belief functions’, particularly by Shafer. For simplicity I will presume that the probability space Ω is finite. Shafer functions can be extended unproblematically to infinite spaces, but the mathematics becomes more complicated, and there are no particularly interesting philosophical issues that arise out of it.

As Ω is a set of possibilities¹⁶, we can define its power set $\mathbf{P}\mathbf{o}(\Omega)$. If we wanted to define a probability function on Ω we could allocate some mass to each possibility. Then the probability of a proposition A on Ω (where A is the union of some possibilities) is the sum of the mass of every element in A . To define a Shafer function on Ω we allocate a mass to every element of $\mathbf{P}\mathbf{o}(\Omega)$. Then $Bel(A)$ in the model is the sum of the masses of every subset of A . The plausibility of A , written $Pl(A)$, is defined as $1 - Bel(\neg A)$. For ease of reference, I'll write Bel_S for a Shafer

¹⁶ I don't call them possible worlds because that would imply there are infinitely many of them. Rather Ω is a finite partition of the possible worlds, with each element of the partition being called a ‘possibility’.

I am calling the power set function $\mathbf{P}\mathbf{o}$, rather than \mathbf{P} as would be standard, so as not to cause confusion with \mathbf{P} which is used to represent a family of probability functions. If I had to I could just differentiate the functions by the different fonts used, but unless it's necessary this seems inappropriate.

function. This way the notation doesn't presuppose that Bel_S functions are reasonable Bel functions.

Formally, when there is a possibility space Ω , a mass function $m: \mathbf{P}\mathbf{o}(\Omega) \rightarrow [0, 1]$ and a Shafer belief function Bel_S the following must hold.

$$\begin{array}{ll} \bullet & m(B) \cdot 1 \\ B \bullet & \bullet \\ & m(B) \cdot Bel_S(A) \\ & B \bullet A \end{array}$$

It should be clear that once the mass function m is defined the Bel_S function is also defined. It is less obvious that once the Bel_S is defined function the mass function is determined, however this can be proven (Shafer 1976: 39). Usually I'll define mass functions only, because they are clearer and shorter.

We can represent prisoner a 's state of belief as a Shafer function. The possibility space Ω is still $\{ab, ac, bc, cb\}$, and the mass function is given by $m(\{bc\}) = m(\{cb\}) = m(\{ab, ac\}) = 1/3$. The Bel_S and Pl functions defined by this mass function are such that for any proposition A on Ω , the prisoner's degree of belief in A , as we have defined it, is vague over the interval $[Bel_S(A), Pl(A)]$.

Shafer (1981) points out that we can think of Shafer functions in the following way. Imagine that we will find out for certain that B , where B is some proposition on Ω . (Shafer uses an analogy with being sent a single coded message.) Assume that for every B we have a precise degree of belief that it will be what we will find. Then $Bel_S(A)$ is our degree of belief that we will find out A . As has been argued (see for example Pearl (1988)) we can think of $Bel_S(A)$ as the probability of A being provable, or more generally of the probability of \mathcal{A} , for some interpretation of \mathcal{A} . In chapter 8 I'll develop a constructivist theory of probability based around this intuition. That theory arguably does a better job than Shafer's own at capturing the intuitions he is promoting.

Intuitively, this seems to show that the Shafer functions can model *some* important epistemic property, but it is doubtful that it is credence. If Bel_S is the appropriate representation for some agent, then she should be convinced that her credence in A is at least $Bel_S(A)$, not that it's exactly $Bel_S(A)$. Even on Shafer's preferred analogy of the secret message, it seems more plausible to say that Bel_S defines a lower bound on credences, not the actual credences. If we can prove A from the message this would be conclusive reason to believe A , but it is possible that we won't be able to prove A yet A still be true. Hence we should say our credence in A is bounded below by our credence in A 's provability. From here I'll assume that the point of a Shafer function is to set bounds on credences, as this seems to be the only plausible interpretation. This move is not original; as will be seen below, there have been some detailed investigations into its plausibility.

Shafer functions are important to this project because of theorem 3.6.1, a proof of which is given in appendix 3B.

Theorem 3.6.1. Let Bel be a vague belief function such that there is a model \mathbf{K}^* for Bel satisfying (1), (3) and (4). Let Γ be any finite field of propositions such that $Bel(A)$ is vague over an interval bounded by rational numbers for all elements of Γ . Then there is a Shafer function Bel_S such that for all A in Γ , $Bel(A)$ is vague over the interval $[Bel_S(A), Pl(A)]$.

The proof works by constructing a mass function from \mathbf{K}^* . The core idea is to allocate mass $1/y$ to the set of possibilities which are left open by p_i for each $i \in \{1, \dots, y\}$. Now to complete the proof that Single Model Approach is stricter than the Many Models Approach, we only need the following theorem, which was proved by Dempster (1968).

Theorem 3.6.2. Let Bel_S be a Shafer function defined on a possibility space Ω . Then there is a convex family \mathcal{P} of probability functions defined on the same possibility space with the following properties.

- (i) For each $A \subseteq \Omega$ and $Pr \in \mathcal{P}$, $Bel_S(A) \leq Pr(A) \leq 1 - Bel_S(\neg A) = Pl(A)$

- (ii) For each $A \subseteq \Omega$, there is a $Pr \in \mathcal{P}$ such that $Pr(A) = Bel_S(A)$.
- (iii) For each $A \subseteq \Omega$, there is a $Pr \in \mathcal{P}$ such that $Pr(A) = Pl(A)$.

Let $W(Bel_S)$ denote the largest such family. Generally, for a given convex family of probability functions, define $inf(A)$ as the lowest value that $Pr(A)$ takes for Pr in the family¹⁷. Then $sup(A) = 1 - inf(\neg A)$ is the maximal value that A takes. It is well-known (see for example Shafer (1976: 5)) that all Shafer functions satisfy the following property, which we'll call **complete monotonicity**. Let A_1, \dots, A_n be propositions (not necessarily disjoint) on Ω .

$$Bel_S(A_1 \vee \dots \vee A_n) \geq \sum_{J \subseteq \{1, \dots, n\}} (-1)^{|J|+1} Bel_S(\bigcap_{i \in J} A_i)$$

For example, for $n = 2$, the condition requires that $Bel_S(A_1 \vee A_2) \geq Bel_S(A_1) + Bel_S(A_2) - Bel_S(A_1 \& A_2)$. However, if we substitute $inf(A)$ for $Bel_S(A)$, complete monotonicity is not a property of all families of probability functions. To take a simple example (this is due to Williams (1978)), consider the family $\{Pr: Pr(A_1) \geq 1/2 \& Pr(A_2) \geq 1/2\}$. In this family, $inf(A_1) = inf(A_2) = inf(A_1 \vee A_2) = 1/2$ and $inf(A_1 \& A_2) = 0$. This shows that the Single Model approach is strictly stronger than the Many Models approach. So if we are to adopt this approach to modelling reasonable imprecise degrees of belief we need a justification for this extra restriction.

All families $W(Bel_S)$ have a further property, which is independent of monotonicity, but which isn't shared by all families of probability functions. I'll call it, for want of a better name, **event defined**. Families which have this property are such that once we have given all the intervals in which the probability of every proposition falls, we have said all there is to say about the family.

¹⁷ Or if this does not exist, the greatest number x such that $x < Pr(A)$ for all Pr in the family. I am generally dealing with closed families here, so I won't attend to the added complications arising from this possibility.

Formally, a family \mathcal{P} is event defined relative to a possibility space Ω iff it meets the following condition.

Event Defined: For all w_i in Ω , let x_i be any element of $[\inf(\{w_i\}), \sup(\{w_i\})]$. Then there is a $Pr \in \mathcal{P}$ such that for all i , $Pr(\{w_i\}) = x_i$.

A family is event defined *simpliciter* iff it is event defined relative to all possibility spaces Ω . Since $\mathcal{W}(Bel_S)$ is simply defined with respect to the interval within which the degree of belief of each proposition falls, it follows it is event defined.

Let Bel be a belief function represented by the convex monotone family of probability functions \mathcal{P} and modelled on \mathbf{K}^* . We will say that Bel is event defined iff \mathcal{P} is. It doesn't follow from (1), (3) and (4) that \mathcal{P} must be event defined. However, if it is not event defined then \mathbf{K}^* is not a good model for it in the following sense. Let \mathcal{P}_E be the smallest event defined family containing \mathcal{P} .¹⁸ The model for \mathcal{P}_E will also be \mathbf{K}^* . Hence the model loses the extra precision of \mathcal{P} over \mathcal{P}_E .

So if we require that all vague belief functions must be capable of being modelled under this approach, we aren't requiring that Bel be event defined, but we are requiring it to be monotone (I will drop the 'completely' unless clarity demands it). What arguments are there, then, that non-monotone belief functions are somehow less reasonable than monotone ones? Well, the plausibility of (1), (3) and (4) is an argument already. However, the non-monotone family of probability functions Williams describes also seems reasonable enough. So we might need some stronger arguments to get to monotonicity.

¹⁸ This family can be constructed in the following way. Assuming we are allowed to quantify over propositions, it is the family $\{Pr: \forall A \inf(A) \leq Pr(A) \leq \sup(A)\}$.

3.6.4 *Uncertainty Aversion*

Schmeidler (1989) has one reason. He argues that any belief function which licences uncertainty averse behaviour will be monotone. I'm not convinced uncertainty aversion is a requirement of rationality. Schmeidler doesn't say it is; I'm adapting his argument somewhat, because some people might reason this way. Schmeidler develops his account of uncertainty aversion by explicit analogy with the standard account of risk aversion. And herein lies the difficulty, for the standard theory of risk aversion is mistaken. That is, there is no way which a person is averse to risk in the everyday sense iff they are *risk averse* in the standard economist's sense. By analogy, Schmeidler's account of uncertainty aversion does not capture correctly the everyday concept of aversion to uncertainty. Hence we have no reason to apply it as a coherence constraint.

Standardly, an agent is said to be risk averse with respect to some good (usually money) iff the marginal utility of that good is declining. More formally, if $U(x)$ is the amount of utility received from x units of that good, the agent is risk averse iff $U''(x) < 0$ for all x . This is thought to represent risk aversion because it means that the agent will always prefer receiving $(a + b) / 2$ units of the good to having a $1/2$ chance of receiving a units and a $1/2$ chance of receiving b units if $a \neq b$. In particular the agent will prefer the status quo to a bet on p , where the chance of p is known to be $1/2$, and she wins a units if p and loses a units if $\neg p$. It is a direct consequence of this definition, and the definition that the expected utility of a gamble is the sum of the utility of each possible outcome times its probability, that if an agent is risk averse they will prefer the mix of any two gambles to at least one of those gambles. If f and g are gambles their mixes, $\alpha f + (1 - \alpha)g$ for $\alpha \in [0, 1]$ are defined as bets which pay whatever f pays if p and whatever g pays if $\neg p$, where p is a proposition whose chance is known to be α .

Following this, Schmeidler defines uncertainty aversion in the following way. The binary relation \geq is a preference ordering on gambles. The set of gambles (what Schmeidler calls 'acts') is called L .

A binary relation \geq on L is said to reveal uncertainty aversion if for any three acts f , g and h in L and any α in $[0, 1]$: If $f \geq h$ and $g \geq h$ then $\alpha f + (1 - \alpha)g \geq h$. Equivalently we may state: If $f \geq g$, then $\alpha f + (1 - \alpha)g \geq g$. (Schmeidler 1989: 582).

I don't want to criticise Schmeidler's extraction of a principle of uncertainty aversion from the classical principle of risk aversion. What I will criticise is that classical principle of risk aversion. The arguments given by Hansson (1988) seem to me to show conclusively that the classical concept of risk aversion does not accurately capture the everyday concept of aversion to risk. Hence we can run similar arguments to show that Schmeidler's concept does not capture the everyday concept of aversion to uncertainty. Any good ideas in what follows here can be found in Hansson.

Joe is a gambler, a person who likes risk. He also is a student of the foundations of mathematics, and as such would like to get a copy of Russell and Whitehead's *Principia Mathematica* (hereafter *PM*). I offer Joe a choice between a copy of *PM* for certain or a gamble. The gamble pays 3 copies of *PM* if a coin lands heads, or nothing if it lands tails. Joe, who has little use for any extra copies of *PM*, chooses the single copy. Hence we can infer that Joe's utility curve in copies of *PM* is convex, hence Joe is risk averse. But we said at the start that Joe liked risk. Indeed we can specify that if I had increased the payout of the gamble to say 6 copies of *PM*, Joe would have accepted it even though he had less use for the 5 extra copies than for the original one. This seems sufficient to say Joe is *attracted* to risk, not averse to it.

The example shows not only that the classical concept of risk aversion is mistaken, but how we might try and fix the difficulty. Let f and g be goods such that the worth to us of g if we already had f is equal to the worth of f . In some cases we can easily measure this. For instance in Joe's case we can ensure that the extra copies of *PM*, whether it's 2 or 5 extra copies, get used the way

Joe would have used them if he already had a copy. So we might give them to Joe's friends and say they're a present from Joe, or if he was to use the surplus volumes to prop up table legs we can give him equivalent sized blocks of wood, and so on. Often we won't be able to make this comparison, but sometimes we will. Now, it seems plausible to say that Joe is risk averse iff he prefers f to a 1 in 2 chance of getting $f+g$. Since we specified that Joe would prefer a 1 in 2 chance of getting $f+g$ to getting f , even when getting $g+f$ was less than twice as useful as getting f , Joe is attracted to risk. Put this way, standard theories of rationality require that everyone be risk neutral.

I conclude that there is no reason to think that the concept picked out by $U''(x) < 0$ is our concept of risk aversion. So the concept picked out by Schmeidler is not our concept of aversion to uncertainty. Hence this can't provide a reason for thinking that it is a requirement of rationality.

3.6.5 Fagin-Halpern Models

Fagin and Halpern (1991) define *probability spaces* as follows. Probability spaces are triples $\langle S, \mathcal{X}, P \rangle$ such that S is a set of possibilities, \mathcal{X} a σ -algebra on S and P a probability function defined on \mathcal{X} . Now for any proposition A , (i.e. for any subset A of S), its inner probability, Pr^* and outer probability Pr^* are defined as follows.

$$Pr^*(A) = \sup \{ Pr(X) \mid X \subseteq A \text{ and } X \in \mathcal{X} \}$$

$$Pr^*(A) = \inf \{ Pr(X) \mid X \supseteq A \text{ and } X \in \mathcal{X} \}$$

It turns out that Pr^* is a DS belief function. More surprisingly, for every Shafer function Bel_S on a set S , there is a probability space $\langle S, \mathcal{X}, P \rangle$ such that $Pr^* = Bel_S$. (A similar relation holds between Pr^* and Pl). The motivation behind Fagin and Halpern's approach is that we can't assign a numerical probability to every subset of S . However, for some sets the agent will have "sufficient information to assign a probability." (349) These are the *measurable* sets. We can work out the

limits on the probability of the other sets by reference to the measurable sets. A similar motivation appears to be at work in Jeffrey (1983a).

Even though everything to this point has been written in terms of bounds, this could be analysed in terms of families of probability functions. Indeed this is exactly the approach Fagin and Halpern take. A probability space $\langle S, \mathcal{X}, P \rangle$ determines a family, \mathcal{P} , of probability functions on S such that for every function $Pr \in \mathcal{P}$ and every element X of \mathcal{X} , $P(X) = Pr(X)$. Then $Pr^*(A)$ is just the minimal value of $P(A)$ in \mathcal{P} .¹⁹

Fagin and Halpern don't appear to take this approach, but we could view their account of probability spaces as an argument for thinking reasonable families of probability functions should have belief functions as their lower bounds. The problem with this argument is that it is hard to see what argument could be given for saying that we should be able to represent vague degrees of belief as probability spaces. A person who has literally no belief might be represented by the vacuous family of probability function $\{Pr: Pr \text{ is a probability function}\}$. For any *a posteriori* proposition, this person's degree of belief in it is vague over $[0, 1]$.

Assume an agent is in this unfortunate state. On Fagin and Halpern's approach, the only way we can move away from this is by coming to have a precise degree of belief in some proposition. If there is no proposition in which the agent has a precise degree of belief our belief, then she must have vacuous beliefs. This seems unreasonable; we can be partially (rather than completely) vague about something without being precise about anything. Hence I don't think it is a requirement of rationality that our degrees of belief be representable by Fagin-Halpern probability

¹⁹ Fagin and Halpern like the analysis in terms of families of probability functions because of concerns about updating. We shall return to these in section 3.8.

spaces, so this can't provide an argument as to why the lower bound on reasonable families of probability functions should be a Shafer function.

3.6.6 Convexity

Even if we were assured that if an agent's representor should, if convex, be monotone, there is a case to be made against requiring convexity in the first place. If this argument appears to succeed it will strengthen our case that the single model approach is too restrictive. Jeffrey (1987) argues that the set of probability functions $\{Pr: Pr(A \& B) = Pr(A) \cdot Pr(B)\}$, which just says that A and B are probabilistically independent, can represent a reasonable epistemic state. However, as can be easily seen, it is not a convex family. The functions $Pr_0, Pr_0(A) = Pr_0(B) = Pr_0(A \& B) = 0$, and $Pr_1, Pr_1(A) = Pr_1(B) = Pr_1(A \& B) = 1$, are both in P , but no linear mixture of them is. Since on the single model approach an agent's representor can only be a convex set of probability functions, the plausibility of Jeffrey's claim that for some A, B this set represents a reasonable state suggests again the single model approach is too restrictive.

3.6.7 Summary

The core aim of this section was to put forward a single model approach to modelling vague degrees of belief. It turned out this approach was stricter than the many models approach of the previous section. Hence I looked at various possible arguments for this extra restriction, but it turned out that none of these were particularly compelling. In the next section I'll look at a different type of model which is less restrictive.

3.7 Many-Urn Models

The models discussed in section 3.6 seemed too restrictive. Belief states which intuitively seem coherent cannot be modelled within the scope of this theory. Hence we might be tempted to look for a different type of model, such as the Many-Urn model set out in this section. I will discuss this model's properties, show how it gets around some of the difficulties of the Single Model

discussed above, and then set out some other benefits and costs of this model. I conclude that it is ultimately not a viable approach, because it is too permissive.

3.7.1 *The Model*

The example of a belief state which couldn't be modelled using the Single Model approach was one represented by the family of probability functions $\{Pr: Pr(A_1) \geq 0.5 \ \& \ Pr(A_2) \geq 0.5\}$. The idea here is that the agent has degree of belief greater than 0.5 in each of two propositions, but is completely ignorant about the probabilistic dependence between these two propositions. We can't represent this in the single model because within that approach there is no way to represent this kind of ignorance. We can have vague degrees of absolute belief, but the allowable vagueness in degrees of conditional belief is somewhat restricted.

Many-Urn models are designed to deal with this specific case. I said earlier it was possible to think of each of the p_i as expressing the proposition that the i 'th ball is drawn from an urn with y balls in it. In Many-Urn models this analogy is developed in the following way. Proposition p_{ij} says that the i 'th ball is drawn from the j 'th urn. It is assumed here that there are several urns, each containing y balls, with a ball to be drawn from each urn. While it is assumed that the initial chance of drawing each ball from a given urn is equal (in a sense the urns are fair) it is not assumed the drawings are independent.

Most, if not all, of the interesting properties of many-urn models are displayed by 2-urn models, so we will restrict our attention to these. We'll start by doing the precise version, and then weaken this to the imprecise version. So p_{i1} is the proposition that the i 'th ball was drawn from urn 1, and p_{i2} the proposition that the i 'th ball was drawn from urn 2. Let P_1 be the set of propositions p_{i1} for all i , and P_2 the set of propositions p_{i2} , again for all i . In this approach, the agents probabilistic beliefs are modelled as propositional beliefs about subsets of $W \times P_1 \times P_2$, with \mathbf{K}^* and Δ^* defined

as above. A model must then satisfy the following two constraints, which are just translations of (1) and (2). The definition of the asterisk operator is trivially extended to the new setting.

$$(1') \quad (S \subseteq P_1 \times P_2 \ \& \ S^* \in \mathbf{K}^*) \rightarrow S = P_1 \times P_2$$

$$(2') \quad \text{If } Bel(A) = x/y \text{ then } \exists S: ((S \subseteq P_1 \vee S \subseteq P_2) \ \& \ (|S| = x \ \& \ S^* \equiv A^* \in \mathbf{K}^*))$$

If a function *Bel* satisfies these conditions, then it must be a probability function. The proof of this is a little less neat, because we can't prove the equivalent of (T6). That is, there doesn't seem to be any contradiction in saying there are sets $S_1 \subseteq P_1$ and $S_2 \subseteq P_2$ such that $|S_1| \neq |S_2|$ and $S_1^* \equiv A \equiv S_2^* \in \mathbf{K}^*$. So we'll just stipulate that *Bel* is a function. Although this possibility isn't formally inconsistent, it is rather odd since we have in some sense $Bel(S_1^*) = Bel(S_2^*)$ and $Bel(S_1^*) \neq Bel(S_2^*)$. Of course this use of '=' isn't formally permissible if *Bel* is a relation rather than a function, but it is a motivation for restricting our attention to belief functions.²⁰

Given this, the proofs that *Bel* must satisfy the three conditions of probability functions follows quite easily. It can be simply observed that *Bel* only takes non-negative values, since sets only have non-negative cardinality, and *Bel* is equal to the ratio of the cardinality of two sets. Similarly the proof that $Bel(T) = 1$ is just as above. The only difficulty is in proving *Bel* satisfies (Pr3). First we'll prove (L4).

$$(L4) \quad \neg \exists i, j \ \forall w: \langle w, p_{i1}, p_{j2} \rangle \notin \Delta^*$$

Assume there are i, j satisfying this condition. Then let $S = \{ \langle p_{\alpha 1}, p_{\beta 2} \rangle : \alpha \neq i \ \& \ \beta \neq j \}$. By the definition of Δ^* it follows that $S^* \in \mathbf{K}^*$. Hence by (1) $S = P_1 \times P_2$, a contradiction. This proves (L4).

²⁰ Formally, the idea I have in mind is that we say $Bel(A) = x/y$ means that the relation $B(A, x/y)$ holds, where the holding of this relation is consistent with the relation $B(A, x_1/y)$, where $x_1 \neq x$.

Now assume that for some A, B such that $\neg(A \& B)$ is provable, \mathbf{K}^* includes $S_1^* \equiv A^*$ and $S_2^* \equiv B^*$ such that $S_1 \subseteq P_1$ and $S_2 \subseteq P_2$. Let p_i be an element of S_1 and p_j an element of S_2 . It follows that for any world w , $\langle w, p_{i1}, p_{j2} \rangle$ is an element of $S_1^* \cap S_2^*$. Since \mathbf{K}^* is closed it contains $\neg(A^* \& B^*)$, so by substitutivity of material equivalents it contains $\neg(S_1^* \& S_2^*)$. Hence for all w , $\langle w, p_{i1}, p_{j2} \rangle \notin \Delta^*$, contradicting (L4). Hence S_1 and S_2 are each subsets of P_1 or of P_2 . By a similar proof we can show that they cannot have any common members, hence $|S_1 \cup S_2| = |S_1| + |S_2|$, and $|S_1 \cup S_2| \subseteq P_1$ or P_2 , so $Bel(A \vee B) = Bel(A) + Bel(B)$, as required.

3.7.2 Imprecision

So moving to a Many-Urn model does not show that Bel , if always rational-valued, may be something other than a probability function. The only interesting difference arises when we move to imprecise models, for there we do perceive a real difference. As above, we drop axiom (2') and import axioms (3') and (4').

(3') Iff $Bel(A) \geq x/y$ then $\exists S: ((S \subseteq P_1 \vee S \subseteq P_2) \& |S| = x) \& (S^* \rightarrow A^* \in \mathbf{K}^*)$

(4') Iff $Bel(A) \leq x/y$ then $\exists S: ((S \subseteq P_1 \vee S \subseteq P_2) \& |S| = x) \& (A^* \rightarrow S^* \in \mathbf{K}^*)$

Again, if $Bel(A)$ always takes precise values (3') and (4') reduce to (2'). However, if it doesn't then we can have a more general system of models. Indeed, we now have a system so general that we can model the family of probability functions we mentioned at the start of this section. \mathbf{K}^* is the following set, with $y = 2$.

$$\mathbf{K}^* = Cn(\{p_{11} \rightarrow A_1, p_{12} \rightarrow A_2\})$$

If Bel can be modelled in this way, then Bel_* and Bel^* have the following rather neat properties.

(M1) $Bel_*(\top) = 1$

(M2) $Bel^*(\perp) = 0$

(M3) If $A \vdash B$ then $Bel_*(B) \leq Bel_*(A)$

(M4) $Bel_*(A) = 1 - Bel^*(A)$

(M5) If $\vdash \neg(A \& B)$ then $Bel_*(A \vee B) \geq Bel_*(A) + Bel_*(B)$ and $Bel^*(A \vee B) \leq Bel^*(A) + Bel^*(B)$

Walley (1991: 600) calls a function that satisfies these conditions 3-coherent. They are regarded as sufficient conditions for a lower probability function by Suppes (1974), Wolfenson and Fine (1982) and Fine (1983). I won't go through the proofs because they are quite similar to the proofs already given. (In particular the proof for (M5) just uses the same move we made to prove (Pr3).) However, this is not sufficient to show that any Bel which can be modelled in this way is coherent. In particular, the following theorem, which seems intuitively plausible, is not validated by this model.

(M6) Assume A_0, A_1, \dots, A_n are pairwise disjoint, and that $Bel_*(A_0 \vee A_1) + \dots + Bel_*(A_0 \vee A_n) > 1$. Then $Bel_*(A_0) > 0$.

If a betting analysis of credences is adopted it is clear why we should want (M6). Assume that $Bel_*(A)$ is our minimum buy-price for A -bets, and that the assumptions in (M6) hold. Then we will be prepared to buy bets on $A_0 \vee A_1$ through to $A_0 \vee A_n$ in such a way that we will have spent more than \$1. This set of bets can only have a return greater than \$1 if A_0 . Hence we are prepared to buy a set of bets which have the same effect as placing a bet on A_0 . But if $Bel_*(A_0) = 0$ then we have said we are not prepared to bet on A_0 at any odds. So there is some incoherence in our betting practices.

Even if we don't adopt a betting analysis there is something wrong with belief functions which don't obey (M6). Assume $Bel_*(A_0) = 0$. Then $Bel(A_0)$ is vague over an interval which includes 0. So it should be coherent with the rest of our beliefs that $Bel(A_0) = 0$. If it isn't we can say we've ruled that possibility out, so we aren't vague over it after all. So assume $Bel(A_0) = 0$. Since for all i in $\{1, \dots, n\}$ $Bel(A_0 \vee A_i) \geq Bel_*(A_0 \vee A_i)$ from the definition of Bel_* , and $Bel(A_0) = 0$ it follows

that $Bel(A_0 \vee A_i) = Bel(A_i)$, so we get $Bel(A_1) + \dots + Bel(A_n) > 1$, a contradiction. So we couldn't have permitted $Bel(A_0) = 0$ after all, so we were not vague over that value.

To show that the Many-Urn model does not validate (M6), assume the possibility space is $\Omega = \{A_0, A_1, A_2, A_3, A_4\}$, and \mathbf{K}^* is defined as the following, with $y = 3$. (This example is derived from one in Walley 1991: 86).

$$\mathbf{K}^* = Cn(\{p_{11} \rightarrow (A_0 \vee A_1), p_{21} \rightarrow (A_0 \vee A_2), p_{12} \rightarrow (A_0 \vee A_3), p_{22} \rightarrow (A_0 \vee A_4)\})$$

For any proposition A , $Bel^*(A)$ is given by the following rules.

If there is no A_i ($i \in \{1, \dots, 4\}$) such that $\{A_0, A_i\} \subseteq A$, $Bel^*(A) = 0$;

If $\{A_0, A_1, A_2\} \subseteq A$ or $\{A_0, A_3, A_4\} \subseteq A$ and $A \neq \Omega$, $Bel^*(A) = 2/3$;

If $A = \Omega$, then $Bel^*(\Omega) = 1$; and

If A satisfies none of the above conditions, $Bel^*(A) = 1/3$.

It can be seen by inspection that \mathbf{K}^* is a model for Bel in the sense set out here. Note particularly that although $(p_{11} \ \& \ p_{12}) \rightarrow A_0 \in \mathbf{K}^*$, there is no subset S of $P_1 \times P_2$ such that $S \rightarrow \perp \in \mathbf{K}^*$ (as is required by (1')), nor is there a subset S of either P_1 or P_2 such that $S \rightarrow A_0 \in \mathbf{K}^*$. Hence $Bel^*(A_0) = 0$. However, $\Sigma Bel^*(A_0 \vee A_i) = 4/3$, contradicting (M6).

3.7.3 Summary

Many-Urn models have a number of nice properties. First, they enable us to remove some of the restrictions of the Single urn models discussed in section 3.6. Secondly, they preserve the conclusion that precise belief functions should be probability functions. Thirdly, the lower bounds of the belief functions they define satisfy the requirements given by Suppes, Wolfenson and Fine for lower probability functions. However, these functions do not satisfy other intuitively plausible

requirements, such as (M6). For this reason I think this approach to modelling probability functions oughtn't be taken.

3.8 Updating

So far I have discussed three possible approaches to modelling vague belief functions. The third of these I found reason to reject, but the other two seemed to be live possibilities. I take it that the failure of the positive arguments in section 3.6 doesn't constitute a negative argument. So this section looks at how credences ought be updated according to these two approaches.

On the many models approach the question is a rather trivial one, as I have already said how each model should be updated. A set of probability functions is updated by updating each function. If an agent's credences are represented by a family of probability functions \mathcal{P} , then their credences after receiving evidence B should be represented by the family \mathcal{P}_B , defined as follows.

$$\mathcal{P}_B = \{Pr(\cdot | B) : Pr \in \mathcal{P}\}.$$

As noted earlier, I take conditional probability as basic, so $Pr(\cdot | B)$ is defined even when $Pr(B)$ is zero. If we adopt a betting analysis of degrees of belief then there is an argument to show this is the only coherent way to update degrees of belief. Walley (1991: 297) uses the betting analysis and a condition he calls full conglomerability to get this conclusion. I don't want to rely on Walley's arguments because, for reasons I'll set out in the next section, I don't find full conglomerability a persuasive restriction on updates.

3.8.1 *Updating Shafer Functions*

The above I take to follow uncontroversially from what I have already shown. However, when we look at the Single Model Approach, and the associated Shafer functions, which updating approach should be adopted becomes a slightly harder question. One of the difficulties is that Dempster and Shafer themselves adopt an updating rule which seems incoherent.

Recall that for every Shafer function Bel_S defined on a possibility space Ω , there is an associated mass function $m: \mathbf{P}\mathbf{o}(\Omega) \rightarrow [0, 1]$. The Dempster and Shafer (DS) updating rule is easiest to explain in terms of this mass function. When we receive evidence B , the possibility space shrinks from Ω to $\Omega \cap B$. The mass of every subset A of this space is given by the following formula.

$$m_B(A) = \frac{m(C)}{1 \cdot Bel(\cdot B)}$$

The idea is that the mass previously assigned to C goes to $C \cap B$. However, the mass that was previously assigned to subsets of $\neg B$ cannot be allocated in this way. So we normalise to take Σm back up to 1. This gives the following formula for the conditional degree of belief in A given B .

$$Bel_S(A | B) = \frac{m(C)}{1 \cdot Bel(\cdot B)}$$

We can define $Pl(A | B)$ in a similar way, or just set it to be $1 - Bel_S(A | B)$. If we try and extend the analysis of updating in section 3.4 to the imprecise models, it seems we will get a similar updating rule. Presumably the updating rule will take \mathbf{K}^* to $Cn(\mathbf{K}^* \cap \{B\})$. More directly, it will take Δ^* to $\Delta^* \cap B$. Now instead of having y dummy propositions, there will only be $y \cdot (1 - Bel(\neg B))$ such propositions left, as $y \cdot Bel(\neg B)$ of them are contained within $\neg B$ in the model \mathbf{K}^* .

Now recall how we defined Shafer functions from the single models. For each p_i , we assigned mass $1/y$ to the smallest set $A_i \subseteq W$ such that $p_i \rightarrow A_i \in \mathbf{K}^*$. Now consider what happens when we take the intersection of \mathbf{K}^* with B . In the new model, the smallest set A_i such that $p_i \rightarrow A_i \in Cn(\mathbf{K}^* \cap \{B\})$ will be $A_i \cap B$. If $A_i \not\vdash \neg B$, this set will be empty. This in effect means that some dummy propositions will be discarded. For the remaining dummy propositions, we will assign mass $1/(y \cdot (1 - Bel(\neg B)))$ to each $A_i \cap B$. This is exactly the same construction as

Dempster and Shafer use to create the updated mass function. Hence viewed this way, when we update a Single Model the updated belief function is determined by the DS updating rule.

3.8.2 *The Three Prisoners Problem Again*

In subsection 3.6.2 we described the Three Prisoners Problem (TPP). At that stage we only discussed the problem of how to represent prisoner a 's initial belief state. Now we want to look at the problem of how to update that prisoner's belief given what the guard says. The situation is entirely symmetric about b and c , so assume without loss of generality the guard says b . Then the only possibilities left are ab and cb . The mass previously allocated to $\{ab, ac\}$ will now be entirely allocated to ab , and then every mass allocation will be scaled up because bc , which previously received mass $1/3$, is now impossible. Hence $m(\{ab\}) = m(\{cb\}) = 1/2$, so $Bel_S(ab \mid ab \vee cb) = 1/2$.

This is absurd. Originally a believes to degree $1/3$ that he will be reprieved. However, if the guard says that b will be executed, his degree of belief that he will be freed rises to $1/2$, and the same happens if the guard says c . But the guard is going to say one of them. As Gardner points out, were this reasoning correct a should believe to degree $1/2$ that he is going to be freed before the guard says anything, since he knows that whatever the guard says that will be his degree of belief.

Formally, what has gone wrong is that our belief rule has breached a condition called finite conglomerability. Let Π be a partition $\{B_1, \dots, B_n\}$ of the possibility space such that for every B_i , $Bel_*(A \mid B_i) > Bel_*(A)$. Then Bel_* is not finitely conglomerable. This seems to be a problem because whatever element of Π is true (and we know exactly one of them is true), $Bel_*(A)$ would rise were we to discover it. We will look in the next section at the plausibility of conglomerability more generally. It suffices to note that here, even without the general rule, it seems the result mandated by the DS updating rule is absurd.

3.8.3 Updating Single Models

If the DS updating rule leads to absurdity, and we update Single Models by using that rule, that model has a difficulty. Fortunately it is one which can be solved. The problem rests with how we interpreted the updated models. In the updated model $Cn(\mathbf{K}^* \cap \{B\})$ there are $y \cdot (1 - Bel(-B))$ dummy propositions. However, not all of these dummy propositions are created equal.

In the precise case, when we updated on B , all of the dummy propositions remaining in the model were on an equal footing. All of the dummy propositions outside B had been discovered to be attached to a false proposition. This is in effect what it is for a dummy proposition to be false. However, we learnt nothing new about whether the dummy propositions inside B were attached to a false proposition. Hence we increase the ‘weight’ of each dummy accordingly. In the imprecise case things aren’t quite so simple. Consider the model for the TPP set out in 3.6.2, which I’ll repeat here.

$$\mathbf{K}^* = Cn(\{p_1 \equiv bc, p_2 \equiv cb, p_3 \equiv ab \vee ac\})$$

When a finds out $ab \vee cb$ (in other words, when the guard says b) he finds out that p_1 is attached to a false proposition. However, while he has found out nothing to discredit p_2 , the new information he has *might* be sufficient to make p_3 false. What went wrong with the approach to updating described above was that we lost this information. All dummies which weren’t definitely ruled out by the new evidence were treated on a par. What it seems should have been done was to divide the dummies into three categories, those that are definitely out, those that are definitely in, and those that might be in or out. The new evidence $ab \vee cb$ puts p_1 into the first of these categories, p_2 into the second and p_3 into the third.

Given this, how do we work out conditional degrees of belief? Or since these degrees will presumably be vague, how do we work out their bounds? To do this we need a more fine-grained taxonomy of what happens to the dummy propositions under conditionalisation. The required

taxonomy depends on the proposition whose conditional degree of belief is being evaluated. We'll call this A .

- (i) α of the dummy propositions definitely go to A .
- (ii) β either go to A or are excluded.
- (iii) χ either go to A or go to $\neg A$ or are excluded.
- (iv) δ either go to $\neg A$ or are excluded.
- (v) ε definitely go to $\neg A$.
- (vi) ϕ are definitely excluded.

By excluded I mean that they are ruled out by the new evidence. In the above example, letting $A = ab$, $\alpha = 0$, $\beta = 1$, $\chi = 0$, $\delta = 0$, $\varepsilon = 1$, $\phi = 1$. The minimal value for A given B will be reached if we exclude all the dummies in (ii) and allocate all the dummies in (iii), (iv) and (v) to $\neg A$. The maximal value will be reached when we assign all the dummies in (i), (ii) and (iii) to A and exclude those in (iv). So we get the following formulae.

$$Bel_*(A | B) = \alpha / (\alpha + \chi + \delta + \varepsilon) \qquad Bel^*(A | B) = (\alpha + \beta + \chi) / (\alpha + \beta + \chi + \varepsilon)$$

Now, it follows from the above definitions that the following formulae hold.

$$\phi / y = Bel_*(\neg B) = 1 - Bel^*(B)$$

$$\alpha / y = Bel_*(A \& B)$$

$$\varepsilon / y = Bel_*(\neg A \& B)$$

$$(\alpha + \beta + \chi) / y = Bel^*(A \& B)$$

$$(\chi + \delta + \varepsilon) / y = Bel^*(\neg A \& B)$$

Combining these two sets of formulae, we get the following.

$$(C1) \quad Bel_*(A | B) = Bel_*(A \& B) / (Bel_*(A \& B) + Bel^*(\neg A \& B))$$

$$(C2) \quad Bel^*(A | B) = Bel^*(A \& B) / (Bel^*(A \& B) + Bel_*(\neg A \& B))$$

These should both have some intuitive force. To make A given B as unlikely as possible, we make $A \& B$ minimally likely and $\neg A \& B$ maximally likely. The reverse applies for maximising the likelihood of A given B . Indeed this result has been reached by a number of different authors. (See for example Fagin and Halpern (1991) and the references therein.) More importantly, as is again pointed out by these authors, if Bel can be modelled by a Single Model and \mathcal{P} is its associated family of probability functions, then the bounds on $Bel(A)$ we get by conditioning \mathcal{P} on B are just those given by (C1) and (C2).

So we have some justifications from within the model to use conditionalisation as our updating rule, even if we don't adopt the Many Models approach. We noted that not every family of probability functions can be represented by a Single Model. In particular only the monotone families can be represented. If we are to adopt the Single Model and update by conditionalisation it becomes important to work out whether the updated family can also be represented. It is rather non-trivial to prove, but Fagin and Halpern (1991) have shown that it always can be. That paper contains their proof and references to some other proofs developed independently of theirs.

3.8.4 *Difficulties*

So far it seems there is nothing wrong with adopting the Single Model as long as we update by conditionalisation and not by DS updating. However, there is a cloud on the horizon. Updating by just using (C1) and (C2), which from now on I'll call 'FH conditionalisation' after Fagin and Halpern, is not commutative. Updating with respect to B_1 and then with respect to B_2 is not necessarily the same as updating with respect to B_2 and then B_1 .

This seems rather odd. After all, it looks like all we've done is conditioned every element of a family of probability functions, and conditionalisation is commutative. The problem arises because the property of being Event Defined is not preserved under conditionalisation. If \mathcal{P} is Shafer bound then when we conditionalise every element of it we will get a new Shafer bound family, and its bounds will be given by (C1) and (C2). However, even if the old family was defined by its bounds, the new family is not. This is somewhat easier to see with an example.

Let A, B, C and D be exclusive and exhaustive possibilities, with the DS belief function on them defined by $m(A) = m(B) = 1/4$, $m(\{C, D\}) = 1/2$. Let $\phi_1 = A \vee B$, $\phi_2 = A \vee B \vee C$. Then FH conditionalisation with respect to ϕ_1 and then with respect to ϕ_2 gives a different result to FH conditionalising with respect to ϕ_2 and then with respect to ϕ_1 . Hence it can't be that each is equivalent to conditionalising with respect to ϕ_1 & ϕ_2 .²¹

The problem can be defined in another way. Recall our definition of $W(Bel_S)$ as the largest family of probability functions bounded by Bel_S . Define $Bel_S(\cdot | B)$ by (C1) and (C2). This is a DS belief function, so we can determine $W(Bel_S(\cdot | B))$. We can also define $W_B(Bel_S)$ as the result of conditionalising every element of $W(Bel_S)$ on B . For any proposition A , the maximal and minimal value of $Pr(A)$ in $W(Bel_S(\cdot | B))$ and $W_B(Bel_S)$ will be the same. However, the following theorem cannot be strengthened to be an equality.

$$W(Bel_S(\cdot | B)) \supseteq W_B(Bel_S)$$

When $W(Bel_S(\cdot | B))$ is strictly larger $W_B(Bel_S)$, we lose information by just modelling the belief state by the boundaries of every proposition. Hence the Single Model approach, even in its most plausible form, has real difficulties. If we insist on modelling agents this way we will lose

²¹ This is pointed out by Paris (1994).

information about their belief states. This loss of information may lead to bizarre results, like the non-commutativity of updating.

3.8.5 *Figures*

These results are not unknown. Fagin and Halpern themselves point out the failure of commutativity under their form of conditionalisation, and supporters of the DS updating rule are aware of its odd consequence in the TPP. Why then are these rules kept? In large part it is because our theoretically preferred model, representing belief states by sets of probability functions and updating them by conditionalisation, has such high computational costs. Given this, some results of van Fraassen (1990) might be important, because they point the way to a reduction in these costs.

Rather than taking sets of probability functions as his given, van Fraassen takes bounds on expected values of gambles. This gives us more information than just looking at the bounds on bets. For example, we saw above that there was no way by looking just at its bounds to represent the distinct information contained in $P = \{Pr: Pr(A_1) = Pr(A_2)\}$. However, we can represent this by looking at bounds on gambles. In this case the gamble that pays 1 if $A_1 \& \neg A_2$, and -1 if $\neg A_1 \& A_2$ will have expectation 0, whereas if P was the set of probability functions its expectation would be vague over the range $[-1, 1]$. We can't do everything we want with just expectation, but as van Fraassen shows we can do quite a bit.

In fact we don't even need to take bounds on gambles as basic. Rather, we can just take as our basic information that certain gambles have non-negative expectation. Every time we say that a gamble has non-negative expectation according to an agent's belief, we rule out some probability functions as being possible members of that agent's representor (i.e. those functions according to which the gamble has negative expectation). Now, say P is the set of all those probability functions such that for each function in P each gamble in a finite class G has non-negative

expectation. If this occurs, say P is a G 's figure. Any family which is a figure of some finite class of gambles is a figure. The complexity of a figure is the size smallest class for which it is a figure.

Van Fraassen gives a constructive proof of the following theorem.

Figure Theorem: If P is a figure of complexity n , then conditionalising P with respect to B results in a figure of complexity at most $n + 2$.

The proof of this is summarised in section 9.2.2. This shows that once we have represented P as the figure of a class of gambles there is a simple algorithm for updating it. Given this, some of the motivation for using incoherent but efficient updating rules should be removed.

3.8.6 Gilboa and Schmeidler's Updating Rule

There are other approaches in the literature to updating vague degrees of belief. The theory of Gilboa and Schmeidler (1993) is based around ν -functions, which are called non-additive probabilities. Again, Ω is a possibility space and propositions are subsets of Ω . There are only three restrictions on ν , which are listed below.

$$\nu(\emptyset) = 0; \nu(\Omega) = 1; \text{ and If } p \vdash q \text{ then } \nu(p) \leq \nu(q).$$

Returning to our prisoner a , the function defined by $\nu(p) = Bel_s(p)$ for the DS belief function defined in subsection 3.6.2. would be a non-additive probability. There are a continuum of updating methods for non-additive probabilities, we'll just consider the extremum cases. These are described by Gilboa and Schmeidler as 'optimistic' and 'pessimistic' rules. The pessimistic rule we've already seen, it is just the Dempster-Shafer rule. The optimistic rule is defined by the following formula.

$$\nu_r(p) = \nu(p \& r) / \nu(r).$$

The TPP seems to be just the right place to apply optimistic and pessimistic rules. The proposition that a is to be reprieved is $ab \vee ac$. Let this be p , and let r be the proposition that the guard says b is to be executed. As we noted above, on the pessimistic view $v_r(p) = 1/2$. On the optimistic view $v_r(p) = 1$. Of course this wouldn't change if the guard had said that c rather than b was to be executed. Optimism is all well and good, but this looks more like fanciful thinking!

We might have been a little unfair to Gilboa and Schmeidler. They formulate their rule in terms of preferences over acts rather than degrees of belief. Still, should a have a choice, then, under the optimistic rule, he should accept a gamble which pays \$1 if he's reprieved and has an arbitrarily high penalty if he isn't. (Of course he already faces such a penalty, so decision theory starts to break down here). Again this seems like optimism gone mad. None of the axioms Gilboa and Schmeidler employ look that implausible, yet their results are absurd. It would be an interesting research program to see where the absurdity creeps in, but this isn't our research program.

3.8.7 Summary

At the start of this section there were two approaches to modelling which seemed plausible. For one of these, the Many Models approach, it is trivial to give an updating rule, because we have already given an updating rule for the precise case. For the Single Model approach, we had to do more work to find an updating rule which was even plausible. However, once we did this we found that updating a belief function can take us from a function which can be accurately modelled (i.e. one that is event defined) to one that cannot. Hence this approach to modelling seems mistaken. The section concluded with a summary of two other discussions of updating vague belief states.

3.9 Conglomerability and Lower Envelopes

Walley (1991: 294-327)²² argues that we should accept as a rationality requirement a principle he calls conglomerability. From this he concludes that (i) probability functions should be countably, not just finitely, additive and (ii) we can't always represent vague epistemic states by sets of probability functions. I think (i) is questionable and (ii) is false, but in each case his argument for these conclusions does not go through. The principle of conglomerability he relies on is so close to principles which are demonstrably not principles of rationality that we shouldn't adopt it as an axiom of rationality. If it falls out from some uncontroversial principles of rationality that our epistemic states should be conglomerative, all well and good. However, arguments designed to show that this principle is intuitively compelling face a fundamental difficulty which Walley does little to overcome.

3.9.1 Conglomerability

In Walley's theory, like many others, the basic conception is the expected value of a gamble. He doesn't require this to be a precise number, so we can introduce vague or imprecise degrees of belief. For example we might be able to say no more than that a certain gamble is more valuable than \$10 and less valuable than \$20. Saying this does not commit us to saying there is some number x in $[10, 20]$ such that the gamble is worth precisely x dollars.

Under fairly minimal assumptions about rationality, *viz* that if gamble α is preferred to β , then $\alpha - c$ will be preferred to $\beta - c$, for some constant c , all the information about an agent is contained in their set of desirable gambles. For our purposes, nothing is lost if we merely talk about greatest lower bounds on the expected value of a gamble α , or what Walley calls lower previsions.²³

²² All page references in this section to this book, unless otherwise stated.

²³ For most of his book Walley works with lower previsions. However, as he notes (160, 616) the class of desirable gambles is slightly more informative, particularly with regard to conditional expectations. Why

What I'll call a propositional gamble is a gamble which takes value 1 at some worlds and 0 at all the rest. These are important in a theory like Walley's for many reasons, not least being that the probability of a proposition is the expectation of a propositional gamble. For any propositional gamble define its associated proposition to be the set of worlds at which it takes value 1. Conversely the gamble which has payout 1 iff a certain proposition is true (and nothing otherwise) is the gamble on that proposition.

For any gamble α , and propositional gamble π , we can define the conditional gamble $\alpha\pi$. As might be clear from the notation, for all worlds w , $\alpha\pi(w) = \alpha(w) \cdot \pi(w)$. So at the worlds where π 's associated proposition is true, $\alpha\pi$ has the same payouts as α , and in all other worlds it has zero payout. So in a natural enough sense this is a conditional gamble. Now we can define the conglomerability axiom.

Let Π be a partition of the possibility space. Every element q_i of the partition has an associated gamble π_i . The principle of conglomerability is that if for all i , $\alpha\pi_i$ is a desirable gamble, then α is a desirable gamble. Walley justifies this principle as follows. (Slight changes to reflect consistency with my notation have been made.)

The conglomerative principle ... requires a type of consistency between current beliefs and current dispositions to update beliefs. It can be justified as follows. Suppose that $\alpha\pi_i$ is desirable for all q_i in Π . This means that You will be willing to accept α after You observe some set in Π , whatever set You observe. Knowing this, You should be prepared to accept α now. (294)

just lower prevision, rather than upper and lower previsions? Well, if x is the lower prevision of $-\alpha$, then $-x$ should be the upper prevision of α . Hence we only need the lower boundaries. The lower boundaries are also easier to compare with other theories, such as Shafer's.

So far we have not mentioned the cardinality of Π . Walley claims that this principle should hold even if Π is an infinite partition. The core of his argument for this is the following line: “Indeed the cardinality of Π seems to be irrelevant in the arguments given to support [this] principle.” (295). If this is right then the principle is in very bad shape, for we’ll see below that it is untenable given a fairly simple assumption when Π is an infinite partition.

3.9.2 Conglomerability and Finite Additivity

Walley notes that if we accept this principle we have to insist on countable additivity of probability functions. The argument turns on the following example, which is fairly well known (321). (See for example de Finetti (1972: 98-100)). Let Ω (the possibility space) be the set of non-zero integers. For ease of exposition, we’ll say x is a random variable ranging over the non-zero integers. Identify Ω with the product space $\Theta \times \Pi$ by identifying each integer n with the pair $(\text{sign}(n), |n|)$. The elements of Π are the pairs $\{-m, m\}$ for all natural m . For simplicity call this the set q_m . Hence the elements of Θ are simply the signs $\{-, +\}$. For simplicity we’ll call the proposition that x is positive $+$, and that it is negative $-$.

Say we have a probability function Pr on Ω with the following properties. For any natural m , $Pr(x = m \mid +) = 2^{-m}$. On the other hand, $Pr(\cdot \mid -)$ is a finitely additive distribution, such that for any natural m , $Pr(x = -m \mid -) = 0$. Finally, $Pr(+)=Pr(-)=1/2$. Obviously Pr itself, as well as its conditionalisation on $-$, is merely finitely additive.

Now let α be a gamble such that $\alpha(x = m) = \varepsilon$ for all positive m , and $\alpha(x = m) = -1$ for all negative m . (We’re assuming here that all bets with positive expected value according to Pr are desirable, and all those with negative expected value are not desirable). Assume ε is an arbitrarily small, but not infinitesimal, positive value. Now α is not a desirable gamble, since its expected payout is $(\varepsilon - 1)/2$. However, for any element q_i of Π , $\alpha\pi_i$ will be desirable. The reason is that $Pr(x =$

$i | q_i) = 1$, and $Pr(x = -i | q_i) = 0$. Hence the expected value of the bet is now $\varepsilon \cdot 2^{-i}$, which is positive. So Pr is not conglomerable.

de Finetti's reaction to this case is rather disconcerting. He shrugs it off as one of those odd results we get when dealing with infinities. Now while we do often get odd results when working with infinite sets, this in itself hardly licences an 'anything goes' approach. We should have a principled reason why this particular odd result is acceptable. If an odd result is forced upon us by a principle which is more compelling than its negation, that result might be acceptable. de Finetti's approach could only work if the intuitive appeal of allowing merely finitely additive probability distributions was greater than the intuitive appeal of conglomerability. And without some more work (part of which is reported below) this can't get off the ground.

3.9.3 *Conglomerability and the Many Models Approach*

My main interest in conglomerability is in a possible argument from conglomerability to an objection to the many models approach. It turns out that if you take conglomerability as your main criterion on an updating rule, and take previsions to be primitive as Walley does, then there are epistemic states which seem reasonable enough but which cannot be represented by sets of probability functions. Or more precisely, they can't be represented by sets of countably additive probability functions, despite being conglomerable. Since for precise probability functions, conglomerability entails countable additivity, this will be a difficulty if conglomerability is a plausible constraint on epistemic states.

This is all consistent with the results of Smith (1961) and Williams (1976) reported above. The state here can be represented by sets of non-countably additive probability functions, but they cannot be represented by conglomerable probability functions. Indeed, there are no conglomerable probability functions which dominate the state in this sense. A probability function Pr dominates a lower prevision \underline{P} iff for all gambles α , the expected value of α according to Pr ,

which I write as $E_{Pr}(\alpha)$, is greater than or equal to $\underline{P}(\alpha)$. Dominating probability functions are important because the way Walley's previsions are represented according to our approach is by the set of probability functions which dominate them.

Before we look at the supposed counterexample, we need one more bit of terminology. We discussed above the vacuous epistemic state, the one represented by the set of all probability functions. In Walley's terminology, the vacuous prevision is where the lower prevision of every gamble is its infimum payout. This wouldn't make sense unless every gamble has a lowest payout, and Walley helps himself to that assumption, apparently to facilitate exposition (58). However, as we'll see below, the assumption does much more work than this.

Here's the example which Walley claims shows the Many Models Approach can't work. Let P_1 be the finitely additive probability function Pr described in the previous section, with Ω , Θ and Π defined the same way. Let P_2 be a finitely additive probability function such that $P_2(x = -m) = 3^{-m}$ for any natural m , $P_2(x > 0) = 1/2$ and $P_2(x = m) = 0$ for any positive m . Now, neither P_1 nor P_2 are conglomerable, however their lower envelope \underline{P} is conglomerable. \underline{P} is defined in the following way. For all gambles α , $\underline{P}(\alpha) = \min(P_1(\alpha), P_2(\alpha))$. The conditional prevision²⁴ $\underline{P}(\cdot | q_i)$, where q_i is any element of Π , is vacuous. To see this, note that $P_1(\alpha\pi_i) = \alpha(x = i)$ and $P_2(\alpha\pi_i) = \alpha(x = -i)$, so $\underline{P}(\alpha\pi_i) = \inf(\alpha\pi_i)$. Indeed, \underline{P} satisfies all of the coherence requirements that Walley discusses.

The next step is to show that no conglomerable probability function dominates \underline{P} . Any such probability function will, by virtue of a more general theorem, be of the form $Pr = \lambda P_1 + (1 - \lambda)P_2$, for $\lambda \in [0, 1]$. Now it is shown that for any value of λ , Pr is not

²⁴ Conditional previsions can be technically defined in the following way. $\underline{P}(\alpha | q) = \underline{P}(\alpha\pi)$, where α is any gamble (previsions are defined primarily on gambles and only derivatively on propositions) and q is the associated proposition of propositional gamble π .

conglomerable. For $\lambda = 0$ the proof is simply the same as for showing the finitely additive function described in the previous section is not conglomerable. For $\lambda > 0$, the proof is slightly more complex.

Note that as $n \rightarrow \infty$, $Pr(x = n) / Pr(x = -n) = 2\lambda^{-1} (1 - \lambda) (2/3)^n \rightarrow 0$, so $Pr(x = n | q_n) \rightarrow 1$. Choose N large enough so that $Pr(x = n | q_n) \geq 1 - \lambda/3$ for $n > N$. Then $P_1(x > -N) = 1/2$ and $Pr(x > -N) = \lambda P_1(x > -N) + (1 - \lambda) P_2(x > -N) \leq \lambda/2 + (1 - \lambda) = 1 - \lambda/2$. However, when $n < N$, $Pr(x > -N | q_n) = 1$, and otherwise $Pr(x > -N | q_n) \geq 1 - \lambda/3$. Hence for all elements q_i of Π , $Pr(x > -N | q_n) > Pr(x > -N)$, so Pr is not conglomerable.

To summarise, the epistemic state represented by \underline{P} satisfies all the coherence requirements that seem plausible. However, if it is a requirement on functions generated by precise models that they be conglomerable, and hence countably additive, that state cannot be represented by a set of probability functions. That is, it can't be modelled by a set of precise models. Hence we should take a different approach to modelling than the Many Models approach, perhaps the approach based on previsions suggested by Walley.

3.9.4 *The Two Envelope Paradox*

There are several problems with the principle of conglomerability. The most serious is derived from some recent work by Broome (1995) and Arntzenius and McCarthy (1997) on the two envelope paradox. It turns out that if we drop Walley's restriction to bounded gambles (i.e. gambles with maximal and minimal payouts) the principle of conglomerability leads to inconsistency. That is, it says of two gambles whose sum is uniformly negative that they are each desirable. To get this result we do not need to postulate gambles with infinite payouts, all payouts will still be finite, but we do need gambles with infinite expected payouts.

As a small aside, it might be noted that these gambles, which have all finite payouts and infinite expected payouts, are rather odd. We can guarantee that the actual payout will be less than the expected payout, which casts some doubt on the principle that we should value gambles according to their expected payout. The argument here will only use the principle that we should value gambles according to their expected payout when that expectation is finite, which seems a more plausible principle.

The two envelope paradox can be set out as follows. One of two envelopes is chosen at random, and an amount of money x is placed in it. The game arbiter calculates x using a chance device such that $Ch(x = 100 \cdot 2^n) = 2^n / 3^{n+1}$ for any integer n , and $Ch(x = k) = 0$ otherwise²⁵. Then $2x$ is placed in the other envelope. The contestant is then given one envelope and asked, before they look inside the envelope, whether they would pay \$10 to swap it for the other one. Assume the contestant knows the method that has been used to calculate the money to be placed in each envelope.

Now at first (and at last I hope) it looks patently absurd that this gamble could be desirable. As far as the contestant knows, there is a complete symmetry between the envelopes. So how could it possibly be worthwhile to trade? Put more forcibly, if the trade is desirable, then presumably it is worthwhile to swap back, again for \$10, since the contestant will now be in the same epistemic position. But the net effect of these two trades is to incur the sure loss of \$20, and that can't be rational. By the normal standards of decision theory this is the worst of all possible mistakes, to incur a sure loss²⁶.

²⁵ A simple chance process would generate this distribution. Start with \$100 in the envelope. Find a fair die tossing mechanism, and keep throwing a 6-sided die until either a 1 or 6 lands face up. After every throw which doesn't land 1 or 6, double the amount of money in the envelope.

²⁶ Scott and Scott (1997) object to the conclusion that, if you'd trade once you'd keep on trading on the

However, look what happens when we bring the principle of conglomerability to work on the matter. Let y be the amount that's in the envelope the contestant is holding, and let Π be a partition $\{q_{100}, q_{200}, \dots\}$ where q_i represents the proposition $y = i$. If $i > 100$, then the expected value of trading given q_i can be easily calculated to be $11i / 10$, and if $i = \$100$ the expected value of trading given q_i is $\$100$. Hence no matter what is in the contestant's envelope, if they knew what it was they'd happily give up that and ten dollars for the other envelope. So by the principle of conglomerability, they should be prepared to give up their envelope and the ten dollars now. That's absurd, so I take it the principle of conglomerability fails here.

Call the envelopes A and B . The core problem, as Arntzenius and McCarthy explicitly point out, is that there are distinct partitions of the possibility space Π_1 and Π_2 such that the expected worth of A is higher (at least $\$20$ higher) given every element of Π_1 and the expected value of B is higher (again at least $\$20$ higher) given every element of Π_2 . Given this it's clear that unrestricted uses of the principle of conglomerability will lead to inconsistency.

Should this be a problem? After all, Walley did explicitly restrict himself to bounded gambles. I think it is a problem. Consider Walley's short argument for extending the principle of conglomerability to infinite partitions. He simply noted that we didn't appear to use the cardinality of the partition in the argument for conglomerability. Well we didn't appear to use the fact that

ground that the reasoning supposedly relies on the contestant in later rounds of trading forgetting the results of their earlier calculation. However, the original argument for trading only goes through if you, in effect, ignore how you would feel were you to be holding the other envelope, so I don't see how this objection works. If you'd ignore the dynamic possibilities of the other position once, there's no reason to not keep on ignoring them.

It might be noticed here that the contestant is making the kind of assumptions which I showed in chapter 2 to be unreasonable. In particular they seem to not be thinking about future trade. But some kind of blindness like this will be necessary if we are to hold a betting analysis of degrees of belief, as all the writers in this field do.

gambles are bounded either, but we must have used one of them, or else we would have ‘proved’ an inconsistent principle. Alternatively, if we really have used neither of these facts, the two envelope case shows that there’s something wrong with Walley’s proof. Unfortunately it is little help in saying what precisely is wrong, but this alone is no reason to hang onto the principle.

Even if it turns out that the principle is consistent when applied only to bounded gambles (and I have no reason to think it is inconsistent in this setting) there remains a philosophical difficulty. We know that when gambles can be unbounded we should not follow the principle of conglomerability. This alone seems enough motivation to deny that the principle is sufficiently compelling to count as an axiom when gambles are bounded. If cases which are roughly alike should be treated in a roughly alike manner, and the principle of conglomerability is provably mistaken when dealing with unbounded gambles, it seems implausible that it could be compelling enough to count as an axiom when we move to bounded gambles. So I conclude that Walley shouldn’t have adopted conglomerability as an axiom. If he did not, he would not have been able to prove that reasonable probability functions should be countably additive, and hence that there are no reasonable probability functions which dominate \underline{P} .

3.9.5 *Intuitions*

The main motivation for the principle of conglomerability is its intuitive plausibility. However, it isn’t at all clear that intuitions in this area should be trusted. Many of the intuitive arguments for conglomerability (or the ‘sure-thing’ principle as it is more commonly known) seem to be equally good arguments for Probabilistic Disjunctive Syllogism (PDS). This is the following argument, where A , B and C are propositions.

$$PDS: \quad Pr(A | B) > e, Pr(A | C) > e, Pr(B \vee C) = 1 \vdash Pr(A) > e.$$

We can run the intuitive argument easily enough. Assume Pr represents your epistemic state. Since $B \vee C$ is known, you can assume that you'll find out one or other of them. In either case your degree of belief in A will be greater than e . So by Reflection, it should be greater than e now. But PDS is clearly invalid as the following example (due to Gillman (1992)) shows.

Assume that a certain deck has only three cards, the ace of hearts, the ace of spades and the two of clubs. A hand of two cards is dealt from it by a fair chance mechanism (i.e. each possible hand is equally likely to be dealt). Let A be the proposition that this hand contains both aces, B be that it contains the ace of spades, and C that it contains the ace of hearts. Clearly enough, $Pr(A | B) = Pr(A | C) = 1/2$, $B \vee C$ is known, and $Pr(A) = 1/3$, violating PDS. The problem is that we moved too quickly between a proposition being true and it being the evidence on which you conditionalise. The following can't both be true (assuming you're reasonable).

- (i) You will find out B only, or you will find out C only.
- (ii) $Pr(A | \text{You find out } B) = Pr(A | \text{You find out } C) = 1/2$.

The intuitive argument for PDS assumed that (i) and (ii) both held in cases like this, which seems to have been the mistake. I only bring this case up to note that intuitions in this area are notoriously unreliable. Many seemingly reasonable people fall for traps like PDS, even though it can be shown with simple cases like this to be a mistake. I suspect that conglomerability *vis a vis* infinite partitions is another principle which, although intuitively plausible at first, is not acceptable on reflection.

3.9.6 Summary

Walley argued that we should adopt conglomerability as a coherence constraint on reasonable models. If we do this there are seemingly reasonable vague epistemic states which cannot be represented by sets of reasonable models, refuting our approach. However, the principle of

conglomerability can easily lead to inconsistency. I argued that this is sufficient reason to be at least open-minded about it in contexts where it is consistent. Hence Walley's refutation does not work.

I might be able to generate the problem again if I accepted countable additivity for other reasons as a constraint on probability functions. In chapter 5 I'll look at some arguments for and against this constraint, with an eye to seeing what problems this will imply for our models.

3.10 Real-Valued Degrees of Belief

It was essential to my proofs earlier that if Bel was precise it should be a probability function that its range was the rationals. This assumption is unduly restrictive and ought be discharged. The most natural way to do this is by defining real-valued degrees of belief in terms of inequalities. Again these definitions are all made relative to some finite field of propositions; I'll take that qualification as implicit in all that follows.

Clearly $Bel(A) = r$ and $r > x/y$ (for integers x, y) should entail $Bel(A) > x/y$. Similarly $Bel(A) = r$ and $r < x/y$ should entail $Bel(A) < x/y$. I'll assume that these entailments hold, and collectively they give a complete account of the meaning of $Bel(A) = r$ for irrational r . As noted above, despite its quantitative representation, we can give a qualitative account of the meaning of $Bel(A) > x/y$. So this theory is again reductive, it says the meaning of the quantitative expression $Bel(A) = r$ can be explained in purely qualitative terms.

Our task now is to work out coherence constraints on real-valued degrees of belief. That is, when are sets of expressions $\{Bel(A) = r_1, Bel(B) = r_2, \dots\}$ coherent. I assume that the set includes an expression for the value of $Bel(A)$ for every A in the field under consideration. We explicate each element of this set as an infinite set of inequalities, and hence the original set is explicated as the union of each of these original sets. The coherence constraint I postulate is just that there be a model for any finite subset of this infinite set of inequalities.

The restrictions on models are just those proposed in earlier sections. Indeed, the models used here are very similar to the ‘single models’ proposed as a representation of imprecision in section 3.6. The idea is that there is a closed set of propositions \mathbf{K}^* on a possibility space $W \times P$ which satisfies the following three conditions:

- (1) $(S \subseteq P \ \& \ S^* \in \mathbf{K}^*) \rightarrow S = P$
- (2) Iff $Bel(A) \geq x/y$ then $\exists S: ((S \subseteq P \ \& \ |S| = x) \ \& \ (S^* \rightarrow A^* \in \mathbf{K}^*))$
- (3) Iff $Bel(A) \leq x/y$ then $\exists S: ((S \subseteq P \ \& \ |S| = x) \ \& \ (A^* \rightarrow S^* \in \mathbf{K}^*))$

Again, I use ‘ \rightarrow ’ for material implication. The number of elements of P is y . From these constraints it can be deduced that Bel must be a probability function. Because the proofs are so similar to those appearing in appendix 3A for the case where Bel takes rational values, I won’t give all of them. Rather I’ll just illustrate how they all work by giving the proof that for disjoint A, B , $Bel(A) + Bel(B) \geq Bel(A \vee B)$. This proof can easily be ‘reversed’ to show $Bel(A) + Bel(B) \leq Bel(A \vee B)$, and hence $Bel(A) + Bel(B) = Bel(A \vee B)$. The proof may be easier to follow after reading appendix 3A where some steps are set out in more detail.

Assume $Bel(A) + Bel(B) < Bel(A \vee B)$. Then there is some integer y satisfying $Bel(A) + Bel(B) < Bel(A \vee B) - 1/3y$. Hence there are integers x_1, x_2 and x_3 such that the following four inequalities hold:

$$Bel(A) < x_1 / y;$$

$$Bel(B) < x_2 / y;$$

$$Bel(A \vee B) > x_3 / y;$$

$$x_1 + x_2 < x_3$$

So by conditions (2) and (3) there are sets S_1, S_2 and S_3 satisfying the following conditions.

$$S_1 \subseteq P \text{ and } |S_1| = x_1 \text{ and } A^* \rightarrow S_1^* \in \mathbf{K}^*.$$

$$S_2 \subseteq P \text{ and } |S_2| = x_2 \text{ and } B^* \rightarrow S_2^* \in \mathbf{K}^*.$$

$$S_3 \subseteq P \text{ and } |S_3| = x_3 \text{ and } S_3^* \rightarrow (A \vee B)^* \in \mathbf{K}^*.$$

Since A and B are disjoint, S_1 and S_2 must be disjoint. But since $x_3 > x_1 + x_2$ there is at least one element of S_3 which is neither an element of S_1 nor S_2 . Call that element p_i . Since $S_3^* \rightarrow (A \vee B)^* \in \mathbf{K}^*$, $\{p_i\}^* \rightarrow (A \vee B)^* \in \mathbf{K}^*$. As the material implication can be contraposed, from $A^* \rightarrow S_1^* \in \mathbf{K}^*$ it can be inferred that $\neg S_1^* \rightarrow \neg A^* \in \mathbf{K}^*$, and similarly $\neg S_2^* \rightarrow \neg B^* \in \mathbf{K}^*$. Now since p_i is both an element of $\neg S_1$ and $\neg S_2$, it follows that $\{p_i\}^* \rightarrow \neg A^* \in \mathbf{K}^*$, and $\{p_i\}^* \rightarrow \neg B^* \in \mathbf{K}^*$. Hence $\{p_i\}^* \rightarrow ((A \vee B) \& \neg A \& \neg B)^* \in \mathbf{K}^*$. Since the consequent is a contradiction, $\neg \{p_i\}^* \in \mathbf{K}^*$, contradicting (1). Hence the assumption must be false. Similar proofs show that Bel must obey all the other axioms of the probability calculus if it is to be coherent.

This style of definition might be incapable of being used to explain what we mean by saying an agent's degree of belief in some proposition takes a particular infinitesimal value. Indeed, I'm not at all sure what we could mean by such an expression. If we can make sense of assuming there is an urn with some non-standard number of balls in it, each of which is equally likely to be chosen, then we can explicate infinitesimal degrees of belief without recourse to a limiting procedure. If, however, we are dubious about the sense of that then recourse to limits won't save the infinitesimals. One way of making sense of these lotteries with infinitely many outcomes might be by moving away from urns and towards natural processes like atom decay or darts being tossed at a dartboard. This is an area that demands more space than it can be afforded here.

Appendix 3A Proof of Theorem 3.3.1

In the proofs I will use S (occasionally subscripted) to refer to a subset of P .

If $\vdash A$ then by closure $A \in \mathbf{K}$, hence $Bel(A) = 1$. Hence (T1). Assume $\neg A$ is a theorem (i.e. $\vdash \neg A$) and $Bel(A) = x/y$. Then there is an S^* such that $|S| = x$ and $S^* \equiv A^*$ is in \mathbf{K}^* . Since \mathbf{K}^* is closed, this

implies $\neg A^*$ and hence $\neg S^*$ are in \mathbf{K}^* . By (L3) this implies S is the null set, i.e. $|S| = 0$. Hence $Bel(A) = 0$, i.e. (T2) holds.

Assume $Bel(A) = x/y$ and $Bel(B) = z/y$. It follows that there are S_1 and S_2 such that $S_1^* \equiv A^*$ and $S_2^* \equiv B^*$ are in \mathbf{K}^* , $|S_1| = x$ and $|S_2| = z$. It follows from (L1) and the way S_1 and S_2 are defined that $S_1^* \vee S_2^* \equiv (S_1^* \cup S_2^*)$ and $S_1^* \& S_2^* \equiv (S_1^* \cap S_2^*)$. We can tell from (L1) that if any two subsets of P , say S_3 and S_4 , are such that $S_3 \equiv S_4$ then $S_3 = S_4$. Finally, if two sets are subsets of P then their intersection and union are also subsets of P . By the transitivity of material equivalences, it follows that $A^* \vee B^* \equiv (S_1^* \cup S_2^*)$, hence $Bel(A \vee B) = |S_1 \cup S_2|/y$. Similarly $A^* \& B^* \equiv (S_1^* \cap S_2^*)$, so $Bel(A \& B) = |S_1 \cap S_2|/y$. We know from set theory that $|S_1| + |S_2| = |S_1 \cup S_2| + |S_1 \cap S_2|$. Dividing both sides of this equation by y and substituting terms which we have already shown to be identical gives us (T3).

If $A \vdash B$ then $A \rightarrow B$ is a theorem. Assume $Bel(A) = x/y$ and $Bel(B) = z/y$. As above, it follows that there are S_1 and S_2 such that $S_1^* \equiv A^*$ and $S_2^* \equiv B^*$ are in \mathbf{K}^* , $|S_1| = x$ and $|S_2| = z$. Since \mathbf{K}^* is closed it contains $A^* \rightarrow B^*$ and hence $S_1^* \rightarrow S_2^*$, or equivalently, $\neg(S_1^* \& \neg S_2^*)$. Assume p_i^* is in S_1^* and not in S_2^* . Since $p_i^* \vdash S_1^* \& \neg S_2^*$, and $\neg(S_1^* \& \neg S_2^*)$ is in \mathbf{K}^* , it follows that $\neg p_i^*$ is in \mathbf{K}^* . By (L1), $\neg p_i \equiv P / \{p_i\}$, so $(P / \{p_i\})^*$ is also in \mathbf{K}^* . However, this contradicts (L2), which says that the only subset Θ of P such that Θ^* is in \mathbf{K}^* is P . Hence there is no such element p_i , so S_1 is a subset of S_2 . Hence $x \leq z$, so $Bel(A) \leq Bel(B)$, as required for (T4).

We have actually proved something more general than (T4). Since all that was used was that \mathbf{K}^* contains $\neg(A^* \& \neg B^*)$, it follows that whenever $A \rightarrow B \in \mathbf{K}$, then $Bel(A) \leq Bel(B)$, provided the agent is rational.

(T5) follows immediately from (T4). Since $A \vdash T$, where T is a tautology, and $Bel(T) = 1$, $Bel(A) \leq 1$. And since $\perp \vdash A$, where \perp is the falsum, and $Bel(\perp) = 0$, $0 \leq Bel(A)$.

From the definitions of degree of belief, it would not be contradictory to say that (T6) failed to obtain. The reason is that $Bel(A) = x/y$ just means that there is some set S of cardinality x such that $Bel(S) = Bel(A)$, not that all sets S satisfying $Bel(S) = Bel(A)$ have this cardinality. This might make us question the use of ‘=’ signs when discussing $Bel(\cdot)$. Fortunately, however, we can prove that it is a requirement of coherence that (T6) holds. Assume that it doesn’t. So there are sets S_1 and S_2 of different cardinality such that the agent believes both $A^* \equiv S_1^*$ and $A^* \equiv S_2^*$ and hence by closure $S_1^* \equiv S_2^*$. Let S_3 be the set $(S_1 \cup S_2) / (S_1 \cap S_2)$. Since S_1 and S_2 are of different cardinality S_3 is non-empty. From $S_1^* \equiv S_2^*$ it follows that the agent believes $\neg S_3^*$, and hence by (L3) that S_3 is empty. This contradicts our assumption, so S_1 and S_2 must be of the same cardinality, as required for (T6).

Appendix 3B Proof of Theorem 3.6.1

I am assuming Bel is such that there is a model satisfying (1), (3) and (4), and trying to prove the lower bound on Bel is a Bel_S function.

$$(1) \quad (S \subseteq P \ \& \ S^* \in \mathbf{K}^*) \rightarrow S = P$$

$$(3) \quad \text{Iff } Bel(A) \geq x/y \text{ then } \exists S: ((S \subseteq P \ \& \ |S| = x) \ \& \ (S^* \rightarrow A^* \in \mathbf{K}^*))$$

$$(4) \quad \text{Iff } Bel(A) \leq x/y \text{ then } \exists S: ((S \subseteq P \ \& \ |S| = x) \ \& \ (A^* \rightarrow S^* \in \mathbf{K}^*))$$

For any $p_i \in P$, let A_i be the strongest proposition such that $p_i^* \rightarrow A_i^* \in \mathbf{K}^*$. That such a proposition exists, and is an element of Γ , is guaranteed by the fact that \mathbf{K}^* is finite and closed under entailments, and that Γ is a field, so if $A \in \Gamma$ and $B \in \Gamma$ then $A \ \& \ B \in \Gamma$. I construct a Bel_S function by constructing its associated mass function m as follows.

Let $S_A = \{p_i: p_i \in P \ \& \ A = A_i\}$. For each proposition A in Γ , $m(A) = |S_A| / y$. Since the sum of the $m(A)$ is 1, m clearly is a mass function and hence generates Bel_S and Pl functions. The only task is to show that the functions thus generated are the bounds on Bel . For an arbitrary A , let $Bel(A) = [x_1 / y, x_2 / y]$ and $Bel_S(A) = x_3 / y$. So the task is to prove that $x_1 = x_3$.

It might be thought that there could be some other cases, such as the possibility that $Pl(A)$ was higher than x_2 . However, this will be dealt with, indirectly, once I have shown that $x_1 = x_3$ for arbitrary A . To prove this I will first prove the following lemma.

Lemma

Assume Bel satisfies the conditions in theorem 3.6.1. Let the bounds of $Bel(A)$ be $[x_1 / y, x_2 / y]$.

Then the bounds of $Bel(\neg A)$ are $[1 - x_2 / y, 1 - x_1 / y]$.

Since $Bel(A) \geq x_1 / y$,²⁷ there is a subset S of P such that $|S| = x_1$ and $S^* \rightarrow A^* \in \mathbf{K}^*$. Hence $\neg A^* \rightarrow \neg S^* \in \mathbf{K}^*$. But $|\neg S| = y - x_1$. So $Bel(\neg A) \leq 1 - x_1 / y$. Assume this bound is not tight, i.e. there is some number $x_3 > x_1$ such that $Bel(\neg A) \leq 1 - x_3 / y$. Then there is a subset S of P such that $|S| = y - x_3$ and $\neg A^* \rightarrow S^* \in \mathbf{K}^*$. But that implies $\neg S^* \rightarrow A^* \in \mathbf{K}^*$, and since $|\neg S| = x_3 > x_1$ this implies $Bel(A) \geq x_3 / y > x_1 / y$. This contradicts the assumption that $Bel(A) \leq x_1 / y$. The other half of the proof, showing the lower bound on $Bel(\neg A)$ is $1 - x_2 / y$ is entirely parallel and hence is omitted.

Now it is a general fact about the Bel_S and Pl functions generated by a common m function that $Bel_S(A) = 1 - Pl(\neg A)$. So if there was a case such that $Pl(A)$ was higher than x_2 , then it would be the case that $Bel_S(\neg A) < Bel(\neg A)$. Hence we need only show $x_1 = x_3$

By definition, x_3 is the number of p_i such that $A_i \subseteq A$. Let the set of all such p_i be S . By the definition of the A_i , this means $S^* \rightarrow A^* \in \mathbf{K}^*$. Hence $x_1 \geq x_3$. Since $Bel(A) = x_1$ there must be some set $T \subseteq P$ such that $T^* \rightarrow A^* \in \mathbf{K}^*$, and $|T| = x_1$. Let p_j be an arbitrary element of T . By the closure of \mathbf{K}^* we know that $p_j^* \rightarrow A^* \in \mathbf{K}^*$, hence $A_j \subseteq A$. So every element of T is, by definition, also an element of S . Hence $x_1 \leq x_3$. From this it follows $x_1 = x_3$ as required.

²⁷ when x, y and z are reals with $y \leq z$, by $x < [y, z]$ I mean $x < y$. Generally when x is a real and S a set of reals, by $x < S$ I mean x is less than every element of S . This usage is quite standard.

Appendix 3C More on Dutch Book Arguments

There are many arguments from the betting analysis of degrees of belief to the conclusion that the probability calculus provides coherence constraints on degrees of belief other than the ones I considered in chapter 2. These include the Savage-style arguments from postulated coherence constraints on preferences, as in Maher (1993) and Kaplan (1996) and the ‘depragmatised’ Dutch Book arguments as in Howson and Urbach (1989), Christensen (1996) and Hellman (1997). All of these as presented make the simple mistake that is fatal to standard Dutch Book arguments of confusing use-value for exchange-value, but whereas this was an incurable problem for standard Dutch Book arguments it might seem curable here. Since these arguments don’t postulate an actual retrade, it might be possible to stipulate, as I did above, that retrade is known to be barred and still get the conclusion that’s desired. Even if this problem can be resolved (and I only accept for the sake of the argument that it can be) there’s a further difficulty at hand; all of the arguments mentioned are question-begging. I have left discussion of these arguments until after chapter 3 because it is easier to see these arguments are question-begging once we know what’s at issue in the dispute. The primary conclusion of these arguments is that degrees of belief functions should obey *Addition*. This is what distinguishes their position from that adopted by, say, Shafer, or indeed most writers who deny that the probability calculus provides coherence constraints. However, *Addition* is a premise of every one of their arguments, which destroys their dialectical effectiveness.

The depragmatised Dutch Book arguments are effectively refuted by Maher (1997), so I won’t say much about them. Maher is rather harsh on orthodox Dutch Book arguments, thinking they are refuted by the declining marginal utility of money. As a rather large class of authors have pointed out (particularly Savage who he is openly following) we can avoid this difficulty by denominating all bets in tickets to a lottery known to be fair. Now this will be ineffective against those who deny that the marginal utility of lottery tickets is constant, but that’s a much smaller class than those who deny *Addition* applies to degrees of belief. The retrade error orthodox Dutch

Book arguments make is sufficient to dispose of them, so this mistake of Maher's is of little consequence.

A typical 'depragmatised' Dutch Book argument is the one in Christensen (1996). He does not believe that an agent who's degree of belief in p is x should be prepared to buy a p -bet for x dollars. However, he does say that the agent should "evaluate such [trades] as fair" (Christensen 1996: 456). So degrees of belief may 'sanction' (his term) certain odds even if the agent does not desire to accept these sanctioned bets. Now making the safe enough assumption that degrees of belief that sanction trades which lead to sure loss are defective he concludes that degrees of belief which do not satisfy the probability calculus are defective. As Maher (1997: 301-3) points out, the argument so far doesn't get the conclusion Christensen wants. Indeed for some simple Shafer functions which are not probability functions no sure-loss trades will be sanctioned. To get Christensen's conclusion, we need the extra premise that if two trades are sanctioned their sum is sanctioned. But given the definition of 'sanction' this just is the premise that degrees of belief must satisfy *Addition*. So the argument is question-begging against the writer who denies *Addition*. Maher shows that similar problems beset the arguments in Howson and Urbach (1989) and Hellman (1997).

The alternative Maher supports is based around 'representation theorems'. A similar approach is taken by Kaplan; I'll focus on Maher simply because the issues that arise are exactly the same. The basic idea is that it is a requirement on the coherence of an agent's preferences that there exist a probability function and a set of utility functions equivalent up to affine transformation such that the agent prefers gamble f to g iff the expected utility of f given the probability function and any of the utility functions is greater than that of g . The probability function will give us the agent's degrees of belief in all propositions. Strictly Maher does not quite believe this, the argument requires that preferences be complete, and Maher does not think this is plausible. If we drop that assumption we get the conclusion that the agent's degrees of belief should be represented

by sets of probability functions, not a single probability function, as has been urged here. However, for convenience he assumes first that completeness holds, and I'll follow this lead. Nothing pertaining to the success or otherwise of the argument turns on this point.

There are a few immediate problems with this approach. Maher needs to assume that if an agent has a higher degree of belief in p than in q they will prefer a p -bet to a q -bet. The retrade problem could arise here, but maybe we can stipulate that retrade is barred and avoid that problem. For the sake of the argument I'll accept that. The bigger problem is that when it is unlikely that we will ever see p or $\neg p$ confirmed we may well prefer a q -bet to a p -bet even if we have a higher degree of belief in p . I would prefer a bet on the Yankees winning the next World Series to a bet on Oswald being Kennedy's assassin, even though I have a higher degree of belief in Oswald's guilt than the Yankees's success, because betting on the Yankees gives me *some* chance of getting a payout. This point is one of the motivations for the intuitionist approach to probability developed in Chapter 8.

Maher is aware of this point, but his attempt to dispose of it is disastrous. His example is comparing a bet on the truth of the theory of evolution, construed as the claim that all life on earth is descended from a few species, with betting on its negation (1993: 89). Taking scientists as his expert function he asks some biologists which of these bets they would prefer, on the assumption that there are extraterrestrials who have been observing earth from its formation and will adjudicate on the bet. He is rather happy that they all plump for betting on Darwin. But this is a perfectly useless result. The objection was that we can have degrees of belief on unverifiable propositions, but our attitudes to bets on these propositions will be quite different to our attitude towards bets on verifiable propositions. He has attempted to counter this by simply making the problematic proposition verifiable. So the realist who thinks meaning and truth conditions go beyond verification conditions will be unsatisfied. And the anti-realist who accepts that verification conditions to determine truth conditions is no happier; she will regard the existence

of the extraterrestrials as new evidence. As it happens the evidence will be both evidence for p and evidence for $\neg p$, but that isn't too strange to the intuitionist. So immediately we have a problem, although again I'd be prepared to accept for the sake of the argument it can be finessed.


The major problem for Maher is that his argument is just as question-begging as the Dutch Book arguments he criticises, though in a more subtle and interesting way. For Maher's argument to work we have to accept some constraints on preferences, such as Transitivity. The argument is only as strong as the argument for these constraints. He has nine axioms which must be justified in some way. The most interesting is Independence, which he construes as follows. \mathbf{D} is in our language a set of gambles and \mathbf{X} a set of propositions, $f \leq g$ means the agent either prefers g to f or is indifferent between them, $f \equiv g$ means that f and g are exactly the same gamble, they have the same payouts in all possible worlds, and $f \equiv g$ on A means that on all possible worlds in which A , f and g have the same payouts.

For all $f, f', g, g' \in \mathbf{D}$ and $A \in \mathbf{X}$, if $f \equiv f'$ on A , $g \equiv g'$ on A , $f \equiv g$ on $\neg A$, $f' \equiv g'$ on $\neg A$ and $f \leq g$ then $f' \leq g'$ (Maher 1993: 190).

The idea is that if f and g have the same payouts given some proposition, say A , varying that payout cannot make us change our preference between f and g . The most famous examples where intuition says this may be violated are the Allais and Ellsburg 'paradoxes'. Since uncertainty plays a larger role in it, I'll briefly sketch the Ellsburg paradox. An urn contains 90 balls. Thirty of these are yellow, and remainder are either black or red in unknown proportion. The payouts for the four gambles in question are given in this table.

	Red	Black	Yellow
f	\$1	0	0
g	0	0	\$1
f'	\$1	\$1	0
g'	0	\$1	\$1

Many subjects prefer g to f , since they know the chance of a yellow ball being drawn but not that of a red ball, but prefer f' to g' since they know the chance of a red or black ball being drawn but not that of a black or yellow ball. With these preferences, and setting A to ‘A black ball is not drawn’ we can see this violates Maher’s independence axiom. No objection yet, many people are just irrational. The real problem arises with Maher’s argument that people who choose in this way are irrational. The following two choice trees set out two ‘tree form’ versions of the choices facing these subjects.



The left-hand tree represents the choice between f and g . The subject is told that if a black ball is drawn they will receive nothing, but if it is not drawn they will have a choice between betting on red and betting on yellow. So far we have a standard enough dynamic choice problem. Maher proposes to make it synchronic by requiring that subjects specify in advance what they would do if they reached the square, that is if a black ball is not drawn. This, he claims, makes the situation

exactly as if the agent was choosing between f and g . Now the right-hand tree is the same as the left-hand tree in all respects but one. If a black ball is drawn the agent receives \$1, not nothing. But the only choice the agent has to make is exactly the same as in the left-hand tree, so they ought make the same choice. We can concede to Maher here that it would be irrational to specify, in advance, a preference for g over f in the left-hand tree and for f' over g in the right-hand tree. This is, however, insufficient for his conclusion.

The problem lies in his assumption that “it seems uncontroversial that the consequences a person values are not changed by representing the options in a tabular or tree form” (Maher 1993: 71). As Seidenfeld (1994) makes clear, this is exactly what is controversial in these circumstances. Indeed this premise, call it Reduction, is expressly denied by a number of heterodox decision theorists, and by writers who deny *Addition* on the occasions they talk about decision theory. There is a good reason for this. Let B , R and Y be the propositions that a black, red and yellow ball respectively is drawn. Many writers will hold that $Bel(B \vee R)$ may be greater than $Bel(B) + Bel(R)$. When evaluating the worth of choosing f' in the original, tabular, it seems plausible that it is $Bel(B \vee R)$ that matters, not $Bel(B) + Bel(R)$. However, in the tree form problem all that matters to f' is $Bel(B)$, for the possibility that we won't need to choose, and $Bel(R)$, for the possibility that we do. Strictly $Bel_{-B}(R) \cdot Bel(-B)$ seems more relevant, but perhaps that can be taken to be $Bel(R)$.

The point is that Maher has to either assume agents only consider $Bel(B)$ and $Bel(R)$ when assessing f' , not $Bel(B \vee R)$, or that $Bel(B \vee R)$ is some function of $Bel(B)$ and $Bel(R)$ so that we can ignore that complication, in his ‘uncontroversial’ assumption. The first option is implausible; surely, when comparing f' and g' , we just compare $Bel(B \vee R)$ with $Bel(B \vee Y)$. More interestingly, I claim that the second is question-begging. Given that virtually everyone agrees that in some cases, for example lotteries, degrees of belief should be probability functions, in some cases the function which gives us $Bel(B \vee R)$ from $Bel(B)$ and $Bel(R)$ must be addition.

Hence he must assume that $Bel(B \vee R) = Bel(B) + Bel(R)$ for the move from tabular to tree form to be plausible. But this is just what he was trying to prove, so the argument is question-begging.

There might be a different argument lurking around here that Maher could fall back upon. He might argue that Reduction is so obvious that if this amounts to *Addition* then *Addition* too is obvious. I suspect this argument has some force with those persuaded by the betting analysis in the first place. So we are back where we started, with the plausibility of Maher's argument standing and falling with the betting analysis. Again this is rather pointless from a dialectical perspective, since most heterodox theorists reject the betting analysis. Indeed, this thesis is nearly unique both in providing a defence of orthodoxy without understanding degrees of belief as dispositions to bet, as I've done in this chapter, and in providing a heterodox theory both motivated and justified by the betting analysis, as I do in chapter 8.

I ought note in advance that when I get to decision theory in Part two I will accept Reduction. However, this assumption plays a quite different role in my argument to that which it plays in Maher's. I already have arguments for *Addition*, so the close connection between Reduction and *Addition* is a reason for me to accept Reduction, not a reason to suspend judgement, as it is here. My decision-theoretic arguments are designed to deduce the consequences of the epistemic principles already developed, not, as in Maher, to develop epistemic principles. So what counts as question-begging is quite different.

Chapter 4

What Probability Is

4.1 The Probability Relation

If probability is to be analysed as reasonable degree of belief, then it must be relational. I have simply assumed this in previous chapters, but there is a rather simple reason for it. What is reasonable to believe depends on what the evidence is. It was reasonable to believe in 1991 that George Bush would win the 1991 Presidential election, but not in 1993 to believe that he did. In probabilistic language, the reasonable degree of belief in *George Bush wins the 1992 Presidential election* for anyone in 1991 was rather high. However, by 1993 the reasonable degree of belief in it became quite low. So consistency demands that probability sentences are also relational in some way. They don't necessarily have to refer directly to the hypothesis whose probability is in question and the evidence. The evidence can be, and often is, referred to by rather circuitous paths. First, it is more common to refer to a rule for determining the evidence than the evidence itself. Secondly, that reference is itself often implicit and determined by context. But a reference to evidence will always be there, as both Keynes and Carnap were careful to point out.

The fact that *probability*₁ [=probability] *is relative to given evidence* and that therefore a complete statement of *probability*₁ must contain a reference to the evidence is very important. Keynes was the first to emphasise this relativity. The omission of any reference to evidence is often harmless if the elliptical nature of the statement was clearly recognised. However, this omission was the general custom with earlier authors, and it often caused lack of clarity (Carnap 1950: 31, italics in original).

Carnap thought he needed to distinguish two types of probability, one of which he refers to in this quote as *probability*₁. This is part of the evidence that probability sentences do contain hidden

references to evidence. Carnap thought that the term ‘probability’ was systematically ambiguous, and generations of theorists have agreed with him. If probability sentences are read as absolute this ambiguity becomes inexplicable, as indeed Carnap thought it was. However, once it is noted that they contain hidden references apparent ambiguity can be explained away as changes in hidden references. But to see this some background is needed.

4.1.1 The Relata Are Propositions

If probability is a relation, what are the relata? Carnap held that they were sentences. I think it better to take the relata to be propositions. There are two reasons for this; basically the two objections that I made to the syntactic nature of Carnap’s theory in section 1.7. The first, which I won’t recount, is that syntax seems little guide to probability in cases where we have doubts about the projectability of the predicates being used. The second, which I said very little about there, is that often we can’t put our evidence in sentential form.

If we take evidence to be sentential, we have to buy into arguments about the theory-dependence of language and about whether or not we can construct an observation language. In particular, it seems that we might be committing ourselves to being on the losing side of these arguments. So I borrow a technique from Lewis (1996) in order to construct a unique proposition from our sense data. Propositions are simply sets of possible worlds. Proposition p is just the set of worlds in which p . For any possible world w , let w be in p iff w is consistent with all our sense data, over our history. We might for some purposes want p to include all worlds which are consistent with our current internal states. This would allow us to, for example, give a non-zero probability to the proposition that the world was created two seconds ago with all our memories hard-wired into us. However, I’ll assume that Lewis’s definition of evidence is generally correct, and that our sensations of yesterday are part of our evidence. Now p is clearly a proposition, though it may not be clear which proposition it is. This may be due to some vagueness about what qualifies as sense data. I am thinking here particularly of phenomena such as blind-sight which might make us think

that it is a vague matter whether a given world is consistent with a subject's sense data or not. But I have avoided the problem of there being no sentence in any particular language which accurately captures all and only our evidence. That such an evidence sentence can't be given counts against taking the *relata* to be sentences; that I can give an evidence proposition counts for taking the *relata* to be propositions.

4.1.2 *Chance*

Given this definition of probability in terms of a relation between two propositions, we can define some related notions, such as objective chance. Various 'events', such as the decay of an atom, are often held to have an objective chance of happening. Even given all the initial conditions we could not in principle say whether or not the atom will decay, just what its chance of decay is. Carnap thought that this required a different probabilistic concept, which he christened probability₂, and analysed in terms of frequencies. Following Lewis (1980), I call the concept 'chance' and claim that it should be understood in terms of propensities, not frequencies. However, unlike Lewis, I think chance can be analysed purely in terms of probability itself, or degrees of reasonable belief. The point of this subsection is to note what this analysis is, and to deal with a technical issue concerning the tensed nature of chance sentences. Defence of the analysis is delayed until section 4.3.

The fundamental fact known about chance is what Lewis calls the Principal Principle (PP). I have implicitly assumed this in previous chapters. The PP says that, if an agent knows that the chance of p at some time is x and has no other 'admissible' information, then the only reasonable degree of belief they can have in p is x . Admissible information is, in Lewis's terms, "information whose impact on credence about outcomes comes entirely by way of credence about the chances of those outcomes." (Lewis 1980: 92) What is admissible is time-dependent. That the suspect fired the gun at midday is admissible information at midday, but not at 11am. For my purposes, 'purely historical' information is always admissible at t . So, at t , if an agent knows the chance of p is x

and knows nothing that is strictly about the future, they have no working crystal balls or such devices, then the only reasonable degree of belief in p is x .

Lewis notes one possible analysis of chance that could be derived from the PP, but for reasons we'll look at in later sections denies that it works. The analysis is derived from how we cash out 'purely historical' information. As is well known, it is very hard to cash out the intuitive notion of the time that a sentence is about from purely syntactic features of the sentence, such as what times are mentioned. Again, a move to propositions saves the day. A proposition is purely historical iff it is true in all possible worlds which are an exact match for this world to the present time. The history of the world proposition for the present is the conjunction of all such propositions. More generally, the history of the world proposition for world w at t is the set of all possible worlds which are an exact match for w to time t . Call this proposition $H_{w,t}$. So as a first approximation, the chance of p at t in w is the probability of p given $H_{w,t}$.

In simpler terms, chance is probability (i.e. reasonable degree of belief) given total history. However, while that is a good enough analysis for 'The chance of p at t is such-and-such', for the apparently simpler sentence 'The chance of p is such-and-such', we need a more complicated analysis. The following approach might be plausible. If the time is now t , the proposition *The chance of p is r* is the set of worlds w such that the probability of p given the history of w to t is r . The problem with this approach is that it makes knowing what time it is a prerequisite for knowing the chance of p . Let w_1 be a possible world such that the probability of p given the history of w_1 to t_1 , which is some time other than t , is r , but which is also such that the probability of p given the history of w_1 to t is not r . Then, on this hypothesis, an agent who knows that the chance of p is r will know that w_1 is not the actual world. This seems a bit strong. An agent who doesn't know whether the time is t or t_1 will presumably not have enough information to know she is not in w_1 .

So we have to complicate the analysis of chance sentences by using a device also due to Lewis, in this case to his (1979b). *The chance of p is r* is a proposition *de nunc*. Put relatively formally, we can say that the possibility space of which this proposition is a subset is not the set of possible worlds but rather the set of tensed possible worlds. A tensed possible world is just a pair of a world and a time. The time can be thought of as ‘the time it is’ in that tensed possible world. Now when an agent knows what time it is, this knowledge doesn’t allow them to say that certain possible worlds are not actual (except perhaps some short-lived worlds). Rather, it allows them to rule out all tensed possible worlds except those where the time is what it actually is. Knowing *The chance of p is r* allows an agent to rule out all tensed possible worlds $\langle w', t' \rangle$ except those where the probability of p given the history of w' to t' is r . In our above example, the agent couldn’t rule out the possibility that the actual *tensed* possible world was $\langle w_1, t_1 \rangle$, and hence couldn’t rule out that w_1 was the actual world, as required.

4.2 Necessitarian Probability

I find internalism about justification highly plausible. It would be odd indeed if what was reasonable for me to believe was unreasonable for a brain in a vat with my sense data, or if what was reasonable for the vat-ensconced brain was unreasonable for me. Obviously, what I know and what the brain knows could be different; the requirement that what is known is true entails that. But the moral of the considerations about what would be reasonable is, I think, that what it is reasonable to believe, what degrees of belief are reasonable in a proposition, is determined entirely by internally available data. Knowing someone’s evidence determines entirely what it is and isn’t reasonable for them to believe.

There is a larger debate about internalism which I could buy into. But that would be a different dissertation. Rather, I’ll let the brain-in-vat considerations flagged in the previous paragraph be enough grounds to take internalism more or less for granted. Now, if we’re internalists, we get a curious result. Internalism says that what it is reasonable to believe supervenes on evidence. This

is a *modal* or interworld supervenience claim. If it is reasonable in world w to believe p to degree d on evidence e , then it must be that it is reasonable in w^* to believe p to degree d on evidence e . No variation in reasonable degree of belief without variation in evidence. So, continuing our example, that a person's entire evidence is e entails that it is reasonable for them to believe p to degree d . And the converse claims also carry through. Had it been unreasonable in w to believe p to degree d on evidence e , that would have meant that a person's entire evidence being e entails that it is unreasonable for that person to believe p to degree d .

To put the same points more dramatically, probability statements are, in their complete form, necessary¹. If the probability of p given q is less than 0.5, then it must be less than 0.5. This might seem like a refutation of our theory. Surely, the objection goes, statements like 'The Yankees will probably win the next World Series' are empirical. Call this sentence S . S is false in the actual world, but we could imagine worlds in which it is true (or conversely if you like the Yankees' chances). The objection misses an important point. Sentences like S are, as Carnap noted, elliptical. S says that the probability of the Yankees winning the next World Series is more than 0.5, *relative to some contextually determined evidence*. I think the most natural is the history proposition of the world to the present, but it might be something else. I'll look at this question in more detail in the next section. Now it is an empirical matter which proposition that is, and with different histories substituted for the elliptical evidence we might get different truth values. But the full sentence, with the elliptical clauses cashed out the way they actually are, S becomes necessary.

¹ Might they even be analytic? Probably not. I argue in section 4.4, that according to a plausible kind of relativism about what is reasonable, probability sentences are necessary *a posteriori*. Even if probability sentences are *a priori*, it requires a rather liberal definition of 'analytic' to count these as analytic.

Many people think that probability could not be a necessary relation (i.e. that probability sentences are either necessarily true or necessarily false), so it is instructive to see how little I have used to get this far. The premises are just three.

1. Probability should be analysed as reasonable degree of belief
2. Internalism about reasonable belief: what it is reasonable to believe supervenes on evidence
3. Whenever all *A*-worlds are *B*-worlds, *A* entails *B*.

Anyone who wants to challenge the notion of necessary probability has to challenge one of these premises or the inference from these to the theory that probability is a necessary relation. And I hope the above has convinced the reader that this would be a non-trivial exercise. None of this should be taken as a recantation of my opposition to logical analyses of probability as in Carnap. We can, and do, say that some degrees of belief are unreasonable given certain evidence without saying there is a distinctive syntactic relation between the hypothesis and the evidence for it. Just because this relation holds in all possible worlds if it holds at all, we aren't required to assume it holds by virtue of syntax.

4.3 Ambiguity and Relations

There is one obvious objection to the theory outlined here. I have said that probability is essentially relational. However, probability sentences seem 'absolute' in some sense. We can say, sensibly, that it is probable that Oswald killed JFK, or that the Yankees will win the next World Series. Here it seems probability is being applied to single propositions, not pairs of propositions. This objection crops up recurrently in the literature, the most recent manifestation being Lowe (1996). He thinks the only way of getting around the problem is to understand these absolute probabilities as probability conditional on a tautology.

If that were the only available move the theory would be in difficulty. However, there are better moves available. I hold that in all cases where no evidence is explicitly mentioned (and indeed in many where it is) there is an implicit reference to evidence. We can regard the sentence as elliptical for a sentence in which the reference to all the evidence is explicit. To make this move we don't need any particularly extravagant linguistic tools. Which evidence is being referred to will usually be determined by what Lewis calls 'conversational score', Quinean principles of charity, and perhaps some general conventions. I'll make some speculations as to the nature of these conventions below. It would be nice to know precisely what the conventions are, but for the defence against this objection I only need the plausible hypothesis that there are such conventions.

To note that the tools I am using are not 'extravagant', we can see that they are needed to explain quantifiers in ordinary language. If we took the domain of any quantifier to be as large as possible, then anyone saying "All glasses are empty" would speak falsely. After all there are always some glasses somewhere in the universe which are not empty. However, this can sometimes be said truly, and the best explanation of that is that the domain of 'all' is implicitly restricted to nearby things, or to things under consideration. Generally the scope of quantifiers, the set of things 'under consideration' will be determined by conversational score. However, sometimes that set will expand because of new items which are drawn to the attention of speakers. Now we can either deal with this by developing more complicated rules for conversational scorekeeping, which is what Lewis does, or we can have a less restrictive context and be more reliant on charity considerations. Either way, whatever tools we need to explain quantifiers will be sufficient to explain the elliptical references that I need, which eliminates, I hope, one potential implausibility.

We can turn this defence against Lowe's objection into an advantage of the theory. Since Carnap (1945), many writers have held that 'probability' is ambiguous, referring sometimes to an epistemic concept (either objective, as here, or subjective, as in some other writers) and sometimes to objective chance. I can explain the data they rely on by means of the elliptical nature of

probability sentences. In the next sub-section I argue that reducing a brute ambiguity, as in Carnap, to a mere ellipticality is a theoretical advantage. In the following sub-section I argue that the ellipticality theory does better than the ambiguity theory at explaining three general features of our usage of probability. Finally I make some speculations about the conventions governing implicit references to evidence.

4.3.1 *Ambiguity and Ellipticality*

I need to make clearly my distinction between ambiguous and elliptical terms. The difference is best brought out by examples. *Bank* is clearly ambiguous in that it might refer either to a financial institution or a riverside. On the other hand, a term like *citizen* is what I'll call elliptical. In various contexts, uttering *Lisa is a citizen* might mean that she's a citizen of Australia, or Britain, or France, or wherever. Which of these is referred to will depend on the conversational score. We can always specify which proposition we were referring to by being explicit about what *citizen* meant, eg by uttering *Lisa is a French citizen*. So, like *bank*, *citizen* can refer to different things on different occasions. However, unlike *bank*, all the possible referents of *citizen* are closely related, they have a large common meaning. Indeed, for some purposes we might say that this common meaning is the meaning of *citizen*. It is clear that there is no particular relation between financial institutions and river sides that grounds their being referred to by the same word. They are not, in the phrase I have used above, different facets of the same concept. So while both *bank* and *citizen* are ambiguous in a weak sense, they both can mean different things in different contexts, in a stronger sense *bank* is ambiguous and *citizen* is merely elliptical.

There are many important terms in philosophy which are elliptical. Definite descriptions like *the dog* will generally be elliptical. Focus on definite descriptions like *the King of France* can make us forget that most of the times we use a definite description we don't assume there is at most one possible referent, but merely a unique contextually relevant referent. Modal terms like *possible* are elliptical; indeed I suspect the ellipticity of *possible* and *probable* is linked. More

contentiously, I suspect that conditionals are elliptical; this would explain the strong connections between different types of conditional (eg indicative and subjunctive) without denying that in a weak sense the conditional form can take different meanings. Even our simplest connective *and* is, I think, elliptical. In one sense it carries Gricean implicatures, and in another it doesn't². But I'm getting away from my main topic.

Why is reducing an ambiguity to an ellipticality an advantage? There are two closely related reasons. The first is a simple appeal to Occam's Razor. If probability is elliptical there is only one distinct 'thing' we have to appeal to in explaining its usage, i.e. its meaning. If it is ambiguous we have to appeal to each of its different meanings. The fewer abstract theoretical entities the better, so the elliptical theory wins. But Occam's Razor is perhaps dubious, particularly when used to argue for quantitative parsimony, not qualitative parsimony.³ Since I am not disposing of meanings generally, just reducing the number of them needed, perhaps the appeal is misplaced.

The second is that the ambiguity theorist has, as I'll stress below, some explaining to do. If the different meanings of *probability* are so distinct, why does English use the same word for them? Indeed, why is this replicated throughout other natural languages? We have invented technical

² To see what I was getting at, compare the following two utterances.

Clark Kent went into the phone booth and Superman came out.

In the weekend football results, Essendon beat Sydney and Brisbane beat Carlton.

The first carries an implication of temporal ordering, but to my mind at least the second does not. It isn't that the implication is cancelled in the second case, it just wasn't there to begin with. If that's right we have to consider the tricky question of what makes maxims applicable.

Even if I'm wrong here, 'and' might be elliptical in another sense. Generally 'and' carries an implication of order. But whether this is temporal order, or order of importance, or some other order will be determined by the context. Given the way I've set up elliptical terms, this is enough to make it elliptical.

The first example above is used by Saul (1997) to make a somewhat different point about implicature.

³ Though Nolan (1997) has argued that quantitative, and not just qualitative, parsimony is a theoretical desideratum.

terms like *chance* and *credence* for the different meanings, and if the natural language usage of these terms was enough like their technical usage, perhaps the ambiguity theorist would have the start of an explanation. However, this is clearly not the case; ‘credible’ is undeniably epistemic, but we can use ‘chance’ in ordinary English to refer to epistemic probability, not what are called in theory objective chances. As I’ll show in the next sub-section, when the ambiguity theorist tries to make this explanation by using new principles, like the PP, the explanatory burden simply moves, so they have to find a plausible explanation for that principle.

What is the sense in which elliptical terms have a single meaning? As noted above, we can say that meaning is the common part between the possible referents. Alternatively, we can say that the meaning for elliptical terms is, like the meaning of indexical terms, just the rule for generating a referent from the context. It is in this sense that I want to say that *probability* has just one meaning, that the probability of p given q is the reasonable degree of belief in p given q . More generally, the probability of p is the reasonable degree of belief in p on evidence q for some contextually specified evidence q .

4.3.2 *Problems for the Ambiguity Thesis*

I claimed above that chance could be reduced to probability given total history. I don’t have a knock-down argument for this reduction. But it seems superior to the competition. In particular it seems superior to the ambiguity thesis: the claim that there are such concepts as epistemic and aleatory probability, and these are ontologically unrelated. As is standard, from now on I’ll refer to these concepts as chance and credence. The credence of p for an agent is, on my theory, the probability of p given the agent’s evidence proposition. The ambiguity thesis is rarely defended, but it is often assumed⁴, which makes it harder to argue against. Such a thesis can’t, I hold, explain the following three facts:

⁴ People who endorse it include Carnap (1945), Russell (1948), Hacking (1975), Shafer (1976) and Gillies

- The one word *probability* is used to refer to both chance and credence.
- Chance and credence have the same calculus.
- Knowledge of chances constrains credences.

Because of the reduction of chance and credence to the single concept I can explain each of these. Ambiguity theorists can, I suspect, explain none of them. Hence there is a strong inference to the best explanation in favour of my position. The rest of this sub-section is devoted to setting out more precisely what each of these three explicanda are, why my approach can explain them and why ambiguity theorists cannot.

The first is in effect already discussed. I only bring it up to note the general rule that the burden of proof falls on those wanting to establish an ambiguity. It could be argued here that it was just ignorance on the part of the populace which has led to this confusion between two concepts being named with the one word. Alternatively an historical explanation might shed light on why these two concepts came to share a name. In this context Hacking's book on the history of probability (1975) might be important. As many ambiguity theorists credit the aleatory / epistemic, or chance / credence, distinction to this book, perhaps something like this is in mind. My explanation for the apparent ambiguity is that I argue *probability* is elliptical, as above. When the elliptical evidence is the agent's evidence proposition *probability* has an epistemic meaning, when it is the history of the world proposition, *probability* has an aleatory meaning.

The second refers to the fact that for most theorists, both chance and credence obey the classical probability calculus. Some ambiguity theorists have a clear defence here, for they believe that chance and credence follow different rules. For example, Shafer (1976) believes that classical

(1991). Gillies believes it is a consequence of taking *probability* to mean whatever it would be most scientifically useful to mean. The influence of Hacking's argument that aleatory and epistemic probability are clearly distinct concepts seems quite widespread in the social sciences.

probability logic is appropriate for chances, but not for credence. He argues that the addition principle, that when A and B are exclusive $Pr(A) + Pr(B) = Pr(A \vee B)$ doesn't hold when 'Pr' is interpreted as credence. So again some ambiguity theorists have an escape. However, those theorists who accept that the calculi for the two are the same have, I think, some explaining to do. That chance and credence should obey the same calculus is explained on my picture by reducing them each to a single concept.

The third refers to the PP, which makes an epistemological link between chance and credence. I can't find a single argument in the literature for why this link should hold on an ambiguity view, even among those who hold onto the ambiguity position and accept the PP. There is an explanation for the PP in Bigelow and Pargetter (1990: 154-159), but it doesn't appear to account for the fact that chance statements are tensed, so the explanation here will be different. In any case, they take it as a consequence of there being an explanation that the different probabilistic concepts we use are, in my terms, different facets of the same concept.

The explanation on my account of the PP is a bit roundabout, but I hope interesting. There is an epistemological principle going back to Locke saying that we should use all the evidence at our disposal. In the terms of my theory of probability, our degree of belief in p should be determined by the probability of p given all our evidence rather than some portion of it. I'll call this Locke's Principle. Occasionally, critics of relational analyses of probability such as the one I'm defending say that no explanation of Locke's Principle can be given by holders of that analysis. This criticism can be found in Ayer (1957)⁵. I will first show how Locke's Principle and my analysis of chance explains the PP, and then show how the analysis of chance can help explain Locke's Principle.

⁵ A somewhat different, but I think also successful, reply to Ayer is found in Chisholm (1989).

By my analysis of chance, if we know that the chance of p is n , then we know that if we had all the possible evidence, ie all the evidence available to the present, we ought believe p to degree n . By Locke's Principle, we ought act on the maximal available evidence set. In other words, it is rationally preferable to have larger evidence sets. The degree of belief which it is rational to entertain on a larger evidence set is preferable to one which is rational to entertain on a smaller evidence set. But when we know the chance we know the degree of belief we would have on the largest possible evidence set, the history of the world to the present. So my analysis of chance and Locke's Principle entail the PP. Hence my analysis can explain the plausibility of the PP. I suspect that an ambiguity theory will have a much harder task making such an explanation.

The analysis of chance also goes some of the way towards making sense of Locke's Principle. When we are considering propositions about the future, e.g. Blue Hands will win the Derby, which we'll call p , we have several competing aims. Arguably, it would be ideal to believe the proposition if it's true and disbelieve it if it's false. Ramsey (1926: 194) for example, says this is the ideal of belief. However, in some circumstances this will be irrational. Even if Blue Hands will go on to win the Derby it would be irrational to be overly confident of this before the Derby is run. Part of the reason for this is that being rational is to have rational habits, and a habit which led to overconfidence in Blue Hands's chances will most likely lead to errors elsewhere.

Aiming to have a true belief about who will win the Derby is not our only aim. We also aim to have justified beliefs and to have our degrees of belief match the chances. The latter requirement is a formalisation of the desiderata noted in the previous paragraph that we be neither overconfident nor underconfident about any particular horse. If we have less than full information, the justified degree of belief in p on our evidence and the chance of p may diverge. However, if we knew everything about the history of the world and we were reasonable in our compilation of that evidence, by definition our degree of belief in p would be the chance that p . The motivation for Locke's Principle is that as we get more evidence, as our evidence more closely approximates

the history of the world proposition, the reasonable degree of belief in p will be more likely to be close to the chance that p . This last step is something of an assumption, so I can't take it as a demonstration that Locke's Principle is worthwhile. However, it seems a plausible enough assumption to at least motivate Locke's Principle, and respond to those critics who argue that on this analysis of probability there is no motivation for collecting more evidence.

This section has argued for, in effect, a reduction of chance and reasonable credence to a single concept of probability. This entails a rejection of the popular ambiguity theory, although I accept that probability is weakly ambiguous or what I call elliptical. The principle argument against the ambiguity theory is that it is forced to say that certain facts which are easily explained under this analysis are giant coincidences, so it is a long way from being the best explanation.

As mentioned above, the analysis of chance defended here is due to Lewis, though he didn't accept it because it has a consequence which he found unacceptable. In response, I'll present a short argument as to why this consequence oughtn't be thought of as unduly odd. The problem is that we have an intuitive idea that there is, in most cases, a range of permissible credences. But if chance is analysed this way, and chances are numerical, it seems that all reasonable belief functions must converge in many circumstances. In particular, whenever any reasonable belief function is conditionalised on a history proposition the output must be a single function. (That is, the result of conditionalising any reasonable belief function on a history proposition will be the same). Now there is a dilemma; do I give up on chances being numerical, or do I give up the intuition that there are a variety of different reasonable degrees of belief. It is never said that we have to give up the idea that there are several reasonable belief functions. (A reasonable belief function for Lewis takes evidence propositions as inputs and has numerical reasonable degrees of belief as outputs). It's just that these disparate functions must have the same result under conditionalisation.

First, note that to the realist about chances, this will be no substantial problem. Such a person will think the history proposition of the world includes facts about chances, so the agent who knows the complete history of the world will know the chance of all events that have chances. Since we aim to have degrees of belief match chances, at least when nothing better is available, there is no reason why such an agent should have any degree of belief in p other than its chance. But the agent was arbitrary, so all such agent's should have the same degree of belief. Lewis is no realist about chances, so this is no good as an *ad hominem*, but it helps frame the debate.

In particular, the qualification in the above, restricting attention to those events that have chances, is important. Anticipating a little the results of chapter 5, it seems plausible that sometimes p has no chance, or perhaps more precisely, the chance of p is not numerical. It is simply an assumption in Lewis that chances are numerical, and without that assumption he doesn't have the conclusion that any two people who know the history proposition must have the same numerical degree of belief in p . Again anticipating a little, I argue that when the probability of p given q is non-numerical, i.e. is vague over $[x, y]$ where $x < y$, then an agent can reasonably believe p to any degree in $[x, y]$.⁶ And when that is allowed, I don't get the really paradoxical results that make Lewis give up this analysis.

On this hypothesis, enough purely historical information would suffice to tell a reasonable believer whether the half-life of radon is 3.825 days or 3.852. What is more: enough purely historical information *about any initial segment of the universe*, however short, would settle the half-life!
(1986: 131, italics in original)

It mightn't be immediately obvious why any short period of historical information would settle the half-life. After all, the half-life of an element of radon now depends on the history of the world

⁶ Or more precisely have a credence in p vague over any sub-interval of $[x, y]$.

until the present. However, each of the possible histories had a certain chance of coming about, so we can calculate the chance of a radon element decaying in a given time as the sum across all possible histories of the chance of that history times the chance of radon decaying in that time given that history. Lewis is assuming a rather strong form of additivity here, surely there are at least uncountably many worlds in question, but I'll let that pass. My point is that there is no justification for saying that the chance of each of these histories is numerical. If Lewis doesn't have this, he won't be able to deduce from initial segments the precise half-life of radon.

If that is unpersuasive, there is another more defensive response to Lewis. The intuition he relies on about the reasonableness of divergence of opinion is developed from our practice in everyday life. Now, in everyday life, no one has knowledge of the complete history of the world. Arguably, in such circumstances, some intuitions about epistemic practices in our world are inapplicable. That is, the cost of accepting a theory in conflict with these intuitions is not as high as it would be were the intuitions about matters with which we have greater acquaintance. Even if one doesn't accept my argument above as to why this is so, Lewis's implicit premise that the intuition is applicable is certainly questionable, and hence he has not knocked out the analysis of chance offered.

4.3.3 *The Elliptical Referents*

I have argued so far that there is an elliptical reference to evidence in every probability sentence. I ought to say something about the content of these references in everyday uses of 'probability'. On my theory this will be an empirical question, so the conclusions of this section are necessarily more speculative than the rest of this dissertation. However, I am confident the answers offered here is at least approximately correct.

My claim is that the implicit evidence in "The probability of p is x " or variants on it, is evidence which is either available to the speaker, or common knowledge, when p is about the past or present

and the total history of the world when p is about the future. There are two rules: future-directed probability sentences are talking about chances and other probability sentences about reasonable credences given publically available evidence. There may be exceptions to these two rules, but in general they are correct.

The original evidence can be determined by looking at the impact of new evidence, particularly new evidence which is conclusive that p . If the new evidence, call it e , shows the original probability sentence was incorrect, then e was in the original evidence. If it shows the original sentence is now redundant or in some way superseded, then e was not. Part of the difficulty here is that telling whether the impact of new evidence is to show the old sentence was incorrect or merely redundant is not straightforward, and in particular is heavily theory-laden.

When p is about the past, my intuition is that someone can speak truly in ignorance of genuinely new evidence that will come out in the future⁷, but not in ignorance of evidence that is, in the context of utterance, taken to be widely known. For example, I think the jurors in the second O. J. Simpson trial, who in effect said that it was more probable than not that Simpson killed his wife, spoke truly. And I don't think it would show they spoke falsely if new evidence completely exonerating Simpson appeared tomorrow. However, I don't think that someone who says the probability that Simpson killed his wife is extremely low (or for that matter extremely high) would be speaking truly, even if all the evidence they had pointed in that direction. One can't make p improbable by deliberately avoiding all evidence pointing to p . Evidence available to the speaker is included in the implicit evidence in order to prevent it ever being proper to say " p but it is improbable that p ".

It is clear from the Simpson case that the relevant evidence cannot be the history of the world. If that were so it would be false to say "The probability that Simpson killed his wife is between x

⁷ And not just by accident.

and y ” unless $x = 0$ or $y = 1$. My intuition is that there are true sentences of that form. So probability sentences about past events are not chance sentences. However, it seems that for future directed sentences the situation is somewhat different. Someone who says “It is probable that Blue Poles will win the 4.15” in between the time that Blue Poles’s trainer has decided to remove him from the race and the time this is publically announced to my mind speaks falsely. What they say is reasonable, but wrong. And a similar intuition carries across to all other future-directed probability sentences. In the earlier case we couldn’t include all the evidence available in theory, the history of the world, because that would make all probabilities equal 0 or 1. Here that problem is not present, so there is nothing to challenge the intuition that the truth of probability sentences, whatever their reasonableness, should be determined by the widest possible evidence set.

4.4 Isms

4.4.1 *Absolutism, Relationism, Relativism and Objective and Subjective*

All through the preceding material the questions of what makes a credence reasonable, and how we know one is reasonable, have been left open. As I showed in chapter 3, sets of precise credences which do not conform to the probability calculus are not reasonable. So, anticipating slightly the results of the next chapter, credence distributions should be evaluated for reasonableness rather than single credences. So the questions become: where do the reasonable distributions come from, and how do we know that they are? To answer these we shall look at the question of whether, and to what extent, the set of reasonable distributions can change over possible worlds. I will consider primarily the question of whether the reasonable distributions are determined by social convention, or whether they are fixed by universal principles. That is, I will be looking to see what extent my theory should be relativist or conventionalist.

First, I’ll define the terms mentioned in the title of this sub-section. Necessitarian theories are occasionally called subjectivist or even relativist simply because what is reasonable is relative to

evidence⁸. If one wants to distinguish the necessitarian approach from one where probability is a physical magnitude, as in Popper's propensity account, this might be useful, but overall it is not a useful notation, for reasons stressed in section 1.8. In part it rests on a misconception that the probability of p is a function of one's evidence. This is not the case. The probability of p given q can differ from the probability of p given r , and if we are speaking loosely, not referring explicitly to our evidence when talking about the probability of p given some evidence, it might seem that if our evidence is different the probability of p is different. But the probability of p given q is independent of whether or not q is our evidence. For clarity, and because terms like 'relativist' and 'subjectivist' have useful meanings, I'll refer to the position that probability must be conditional as relationism. If I add to this the position, implicit in what's above, that the set of reasonable distributions is the same in all possible situations, I'll call the position absolutism. Relationism is consistent with all the other positions I'll discuss; indeed, they all imply relationism.

Relativism and conventionalism both say that the set of reasonable distributions is determined in some way by social conventions. Conventionalism says that it is entirely a matter of conventions what is reasonable. These conventions are mostly implicit, and in part may be determined by entrenchment rules in the way Goodman says. Part of the reason that we can infer that probably all emeralds are green from the evidence we have is because *emeralds* and *green* are well entrenched. And entrenchment rules determine conventions of reasonableness. Relativism is like conventionalism in saying that the set of reasonable measures is determined entirely by social conventions, such as entrenchment rules. However, according to relativism, the conventions rigidly determine the set of reasonable distributions. So on a relativist approach, had the conventions been different this wouldn't have affected the probability relation between any two propositions. Because relativism uses rigid designation, if we adopt this theory then there will be

⁸ For example in Popper's (1933) he refers to Keynes's theory as a subjectivist theory of probability.

some contingent *a priori* truths. In particular, it will be true *a priori*, but only contingently, that following social conventions as to what is reasonable, as to what constitute reasonable distributions, is reasonable.

The only difference between relativism and conventionalism is in how they deal with counterfactuals. And I think relativism deals with them somewhat more appropriately. It sounds very implausible to say that what Sally is doing is rational but that it would have been irrational had the conventions been different. It isn't clear that conventionalism is consistent with the internalism I have assumed from the start. If I am being conventionalist then I have to say that the society relative to which conventions are determined is so small that it impossible to have people in different worlds with the same history but different social conventions, else I would have allowed that people with the same histories could have different reasonable degrees of belief.

So the only plausible construal of conventionalism is where the conventions are set in such a way that the agent has evidential access to them. Even if I allow this, the conventionalist is still in some difficulty, for reasons set out in section 1.7. The basic problem is that conventionalism assumes agents have privileged access to what the conventions in their community happen to be. It turns out to be always unreasonable to not believe that the conventions are what they happen to be. But this is just an empirical matter, so agents should be entitled to make reasonable mistakes about this, as about everything else. So the conventionalist position is implausible.

Relativism is untouched by these attacks, and there are some positive arguments for relativism. The main one has already been mentioned. The absolutist is left relying on implausible notions like Moorean intuition to ground reasonable beliefs. She can avoid this by saying evidence about conventions ought to affect our judgements about other matters, but only at the cost of denying some intuitions about what counts as relevant evidence. Alternatively, the absolutist can say that our actual practice of making probability judgements is based on convention, but this is not

necessarily optimally rational. Other communities could show us we have been too restrictive or too permissive in constructing the set of reasonable measures. Presented in this way, absolutism absolves itself of the initial implausibility associated with saying that in a strong sense we have latched on to the one true probability logic just through the ‘swamp metaphysics’ of our language. In conclusion then, I take it that both relativism and absolutism are live possibilities, but conventionalism is not.

Even if we were to adopt relativism, I don’t think we should stop describing the analysis as an objective theory of probability. Carnap says his theory is an objective theory because, “if a certain probability₁ value holds for a certain hypothesis with respect to a certain evidence, then this value is entirely independent of what any person may happen to think” (1950, 43). O’Donnell (1989) in his discussion of Keynes makes a similar point. Now whether this holds on a relativist view depends on what we take him to mean by ‘independent’. At its most natural it means that had people thought differently about what is reasonable, this would make no difference to probability values. And that is a claim which the relativist can endorse. I think the most natural thing for the relativist to say is that probability is an objective relation, but which objective relation it is is determined by what conventions are prevalent in the actual world.

4.4.2 *Contextualism and Belief*

In his (1996) Lewis advocates a contextualist account of knowledge. I want to extend that account here to reasonable degrees of belief. On Lewis’s picture, an agent knows that p in a certain context iff all the contextually relevant possible worlds consistent with the agent’s evidence are p -worlds. Probability functions have as their domain a set of subsets of the possible worlds. As Lewis stresses, in practice ‘the possible worlds’ should not be taken to include all the possible worlds there are, rather just the relevant ones. Propositions all of whose elements are irrelevant will simply not be assigned a probability, or in effect receive probability zero. In this way, context can affect what counts as a reasonable probability function. If the context is such that p is not a

relevant possibility it might be reasonable to, in effect, assign p probability zero, even if this would not be reasonable were the context changed so that p were relevant⁹.

The contextualist turn makes it possible to defend a theory which has come in for some heavy criticism recently. With all this talk about degrees of belief, the reader might be wondering what has become of our familiar absolute concept of belief¹⁰. One rather obvious answer is to say that an agent believes p just when their degree of belief in p is one¹¹. The objection to this from devotees of the betting analysis is to say that we can believe p even when we wouldn't be prepared to bet on p at any (finite) odds. This objection is run, for example, by Maher (1993) and Kaplan (1996) who both define belief in terms of cognitive utilities¹². The problem with this objection is that it equivocates over context. One of the ways that a world can become relevant for Lewis is that we consider the possibility that it is actual. Now when we are deciding whether to bet on p at extremely short odds, it seems pertinent to consider what would happen if $\neg p$. Hence at least one $\neg p$ world is relevant. If there were no $\neg p$ worlds previously relevant then on all reasonable probability distributions the probability of p is 1, because probability distributions are normalised measures across relevant worlds, hence the agent's degree of belief in p might be 1 until the idea of the bet is raised, at which stage context changes, and perhaps what degrees of belief in p are reasonable change. This is related to the possibility of finkish dispositions to bet mentioned earlier that plague crude betting analyses of degrees of belief. In sum, on the contextualist story it is

⁹ As stated this is trivially the case when p is a proposition which is about the context. It is not intended in this trivial sense, as I'm sure most readers will have understood.

¹⁰ And for that matter of 'knowledge', but that's too long a story to enter into here.

¹¹ Although I will use Lewis's work to defend this answer it is not one that to my knowledge Lewis endorses.

¹² They each accept the bizarre conclusion that an agent can believe p when their degree of belief in p is less than $1/2$. Given the broadly functionalist methodology driving the betting analysis, I can't see how they can say an epistemic state such that the agent prefers a $\neg p$ -bet to a p -bet could possibly play the functional role of a belief that p and hence be a belief that p .

consistent that an agent would believe p to degree 1 yet not accept a bet on p at very short odds were it to be offered, hence Maher's and Kaplan's first objection fails.

That objection was intended to show that believing p to degree 1 is not necessary to believe p . It has also been argued that it is not sufficient. Let t be the length of time until this tritium atom decays. For all $x \in \mathbb{R}^+$, $Ch(x = t) = 0$, so for all x an agent who knows the chances involved will believe t is not x . But this is nonsense: if beliefs must be closed under infinite conjunction this amounts to the belief that the atom will never decay, and even if they need not be closed under infinite conjunction, the particular beliefs seem implausible. Lewis's response¹³ is to say that this argument makes a rounding error. The chance that $x = t$ is not zero, rather it is a positive infinitesimal. Hence the agent's degree of belief in $x = t$ should be positive, and hence they need not believe that t is not x . If we can accept infinitesimal degrees of belief, this will be a perfectly acceptable answer. However, for the reasons mentioned at the end of chapter three, I'm not convinced these are acceptable.

So my conclusion here is somewhat hesitant. If infinitesimal degrees of belief are acceptable then belief just means degree of belief 1. If they are not then an agent believes p just in case p is true at all points in the possibility space over which the probability functions in their representor are defined. In that case believing p to degree 1 will be necessary but not sufficient for belief. The only objection to this makes two controversial steps. The first is to infer a disposition to bet from the agent's degree of belief. The second is inferring a conditional about the agents actions from that disposition. This amounts to assuming the disposition is not finkish. However, we have good reasons for thinking that if the agent has the disposition it will be finkish, so this objection fails.

¹³ In his (1994).

4.5 Lewis's New Principle

The story about chances above relies heavily on the Principal Principle (PP). Recently, however, Lewis has argued that the PP should be modified to take account of what he calls 'undermining'. The PP is replaced with the descriptively entitled 'New Principle' (NP). In this section I will argue that the existence of undermining should be more of a problem for particular views of chance than for the PP, and that there are independent reasons for thinking views of chance that permit undermining are flawed.

To motivate the differences between the principles we need to say a little more about the role of 'admissible' evidence in the PP. Obviously, just knowing that the chance of p at some time t is x doesn't entail that the only reasonable degree of belief in p is x . After all, if t is ten minutes ago, and you now know that p , the reasonable degree of belief in p is 1, not x . So we have to restrict what else can be known before the PP is applicable. Lewis doesn't do this completely, but speculates that two types of evidence are 'admissible'. The first, which I have spent some time on, are history propositions. The second are what Lewis calls 'theories of chance'.

A theory of chance is a set of history to chance conditionals. A history to chance conditional has, as the name implies, a history proposition as antecedent, and a proposition about chances at the end of that history as consequent. They are admissible because their impact on degrees of belief goes entirely via their impact on beliefs about chances. Lewis thinks that history to chance conditionals are nomic. Given this, it is worth recalling the theory of laws that Lewis wants to defend.

Among the true sentences, some are quite simple. For example, "Socrates is wise" is about as simple a sentence as one could get. Some are quite informative, like the conjunction of all the sentences in a history book¹⁴. And some manage to do quite well on each of these counts. Those

¹⁴ Assuming the book miraculously contains only truths.

all men know as the laws. That's the spirit of Lewis's theory, but three complications are needed to get it right. First, we assess simplicity and strength (informational content) on a system-by-system basis, not a sentence-by-sentence basis. Secondly, not just any set with a sufficiently high reading on informational content and simplicity will do, only the best such set counts as a law. However, since Lewis thinks in practice the winner of this contest will be so far ahead of the field that debates about how to trade-off simplicity for strength will be moot, this point has little practical importance. Finally, simplicity has to be relative to a language; in a language with no predicate for wisdom "Socrates is wise" will not be at all simple. So we stipulate that we are concerned with the language in which the predicates are the "real" universals.

These predicates, and the laws which they make up, do quite a deal of work for Lewis. "If you're prepared to agree that theorems of the best system are rightly called laws, presumably you'll also want to say that they underlie causal explanations; that they support counterfactuals; that they are not mere coincidences; that they and their consequences are in some good sense necessary; and that they may be confirmed by their instances." (Lewis 1994: 478-9)¹⁵ As I mentioned in the discussion of Carnap in section 1.7, I am rather sceptical that any such set of predicates which is wide enough to support all the ampliative inferences we make (particularly in social sciences) can be narrow enough to avoid licensing inferences which are unsound, but for now I'll let that slide.

The difficulty Lewis sees with this picture is that laws, being 'in some good sense necessary' must have no chance of being false. But laws, and hence chance statements, don't supervene on history: the relevant supervenience is on the complete facts about the world past, present and future. This opens up the possibility of undermining. Formally, let T be the theory of chance for this world. Since T is either a law or the consequence of some laws, the chance of T must be 1. As chances

¹⁵ All references in this section, unless otherwise indicated, are to Lewis (1994).

don't supervene on history, there may be worlds which have the same history as this one, but in which T is not the law.

Let F be some particular one of these alternative futures: one that determined different present chances than the actual future does. F will not come about, since it differs from the actual future. But there is some present chance of F . That is, there is some present chance that events would go in such a way as to complete a chancemaking pattern that would make the present chances different from what they actually are.

The present chances *undermine* themselves (482).

Undermining is certainly peculiar. How it works in practice can be illustrated by taking a simple frequentist account of chance¹⁶. There is some positive chance that every tritium atom from now to the end of the universe will decay in under four days, despite tritium having a half-life of 12.26 years. The chance may be infinitesimal, but Lewis allows that. Indeed he invokes enough infinitesimals to revert back to an idea of Keynes's, saying that the chance of an event at t is zero iff the laws and facts at t rule out that event, and the reasonable degree of belief in p is zero iff p is inconsistent with the evidence available. Now if this strange pattern of tritium decays were to happen, the chance of a given atom decaying in under four days now would be higher than it actually is, because the frequency of decay inside four days would be massively higher.

This makes undermining look at least possibly inconsistent. Lewis invokes the PP to get a formal inconsistency. Let E be any proposition satisfying two conditions. First, it specifies the chance of F , the proposition mentioned in Lewis's quote above. Secondly, it contains no inadmissible

¹⁶ This illustration is used by Lewis.

information about the future; “it does not give any information about how chance events in the present and future will turn out” (483). From the PP we get the equation:

$$Pr(F | E) = \text{Chance}(F)$$

Now take a particular case, let E be the whole truth about present chances, and F be as above. Since F has some chance of coming about $Pr(F | E) > 0$. But F and E are inconsistent, if F is true then chances are different to what they actually are, i.e. E is false. So $Pr(F | E) = 0$. Contradiction.

As is common, a distinctive feature of a theory of probability can be mirrored in a non-probabilistic theory. Usually this mirroring will be between a writer's probabilistic and absolute epistemology. Here the mirroring is between Lewis's probabilistic and absolute metaphysics. Whether or not L is, and always has been, a law is not determined by the present and past ‘history of the world’, or the Humean facts to the present. So there may be a future, say G , such that L and all other actual laws are true in G but a world with our past and present and G as a future does not have L as a law. So is G possible in the ‘good sense’ Lewis mentions? Well it is consistent with the laws, so yes. But if G were true, the past would be changed, L would no longer be a law then as it actually was. And it is impossible in this sense to chance the past¹⁷, so G is impossible. Is this a contradiction? Perhaps not; perhaps there is an equivocation on ‘possible’ here, as frequently happens in discussions of possibility and time. Still, it seems that laws can undermine themselves just as effectively as chances.

What ought our response be to the possibility of undermining? Lewis responds by qualifying the ‘admissibility’ criteria in the PP in two ways. The first is to say admissibility admits of degree. Some information may be strictly inadmissible, so knowledge of it waives our obligation to correlate degrees of belief with known chances, however very nearly admissible, so reasonable

¹⁷ As was stressed by Lewis (1976a).

degrees of belief must still be close to known chances. Secondly, and self-evidently, the admissibility of information may be different for different propositions. A may contain much evidence relevant to whether or not p which does not impact on beliefs about the chance of p , but little information of this kind about q . With these tools in place we can say that information about present chances is for almost all propositions almost entirely admissible. This resolves the contradiction that Lewis had uncovered in his earlier position. The propositions E and F he used were inconsistent, so E is completely inadmissible relative to F . However, relative to most propositions E will be almost entirely admissible.

So the PP cannot be used to generate inconsistent requirements on rational agents. However, Lewis still worries that it can be used to generate incorrect requirements. So we get a slight modification of the PP. Letting H be the history proposition to the present, T the theory of chance for the world, and A any proposition, the old principle said that:

$$\text{Chance}(A) = \text{Bel}(A \mid HT).$$

We were led into difficulties when A and T were too closely related. The new principle fixes this by conditionalising the left hand side as follows:

$$\text{Chance}(A \mid T) = \text{Bel}(A \mid HT).$$

This reflects the imperfect admissibility of chance sentences. By such moves, Lewis rescues enough of his view of chance to be satisfied. The more obvious response to undermining is to see it as a problem for Lewis's view of chance. His account makes undermining possible, but it oughtn't be possible, so we ought reject Lewis's view. There are two reasons we might wish to regard undermining as impossible. The first is a related set of intuitions about chance, the second a plausible principle about chance which is incompatible with the existence of undermining.

The intuitions that are incompatible with undermining revolve around the tensed nature of chance statements. At some level, a statement that the chance of p at t is x is 'about' t . This could be cashed out in many ways. We might say that whatever happens beyond t is irrelevant to the truth of the statement, that it is possible in theory to know whether or not that statement is true at t , or that its truth value cannot be changed by events that happen after t . Any of these renditions will make it undermining impossible. On Lewis's view of chance, there is no sense in which a chance statement is about t . As already noted, it may still be contingent in some good sense¹⁸ at t whether or not the sentence is true, because it might be undetermined whether or not the chance statement will be undermined. So we cannot know that chance statements are true, since to know these requires knowledge of the laws, and knowledge of those requires unattainable knowledge of what will happen in the future.

We can put the same point in terms of the powers of agents. Tim can't kill his grandfather, because whether or not Grandfather died at a certain time in the past is invariant on whatever subsequently happens in either internal or external time. However, on Lewis's view Tim can affect what the chance was of grandfather dying at a past time. What that chance was is not invariant on what happens in later external time.

The objection here mirrors the objection raised by Peter Menzies to Lewis's analysis of causation. As he does with chance, Lewis analyses causation in 'global' terms. Whether or not A causes B in a given world depends not just on 'local' facts about A and B and their immediate environment, but on global facts which determine what the laws are for that world. Lewis needs arguments to defeat the intuitions that private causation and private chances are possible.

The second objection to undermining turns on a principle first enunciated by Bigelow, Collins and Pargetter (1993), something they call the Basic Chance Principle (BCP). The BCP says that

¹⁸ Whether that sense is nomic contingency will turn on definitions of nomic that are at issue here.

if the chance of p in world w at time t is x , and $x > 0$, then there is some world with the same history as w to t in which p is true and the chance of p at t is x . The idea behind it is that in saying the chance of p is greater than zero, we are saying p is in some sense now possible. That is, we are saying in some world like this one, i.e. in which the history is the same and the chance of p is the same, p occurs. It follows from the BCP that chances must supervene on history, and hence undermining is impossible. If we're committed enough to theories which insist undermining is possible we may have justification for dismissing the BCP as an excessively onerous burden on theory, which is what Lewis does. I bring it up here to note there is a general reason for insisting on supervenience of chances on history, beyond the intuitions about localness mentioned above.

These two arguments point against views of chance which permit undermining, and in favour of views which make current chances supervene on history to the present. As argued earlier in the chapter, some of Lewis's objections to such a view turn on the unjustified assumption that chances are numerical. Once that assumption is dropped, the analysis of chance as 'objectified credence' I advocated seems to do all the work we could want. And that analysis is compatible with, indeed entails, the old PP, so I have no need to move to the new principle.

Chapter 5

Supervaluations

5.1 Introduction

When we say the probability of p is x , we mean, roughly, the reasonable degree of belief in p is x . If we could guarantee that for any evidence there was a unique degree to which it was reasonable to believe p , that would be the end of the story. This is the path that Carnap followed, but as I showed in earlier chapters there are several reasons to think it will not lead to a happy ending. Ordinary credences can be vague or imprecise, and there often seems to be different credences which are reasonable on a given body of evidence.

So I have to be more careful when stating the relationship between probability and reasonable degree of belief. The idea defended in this dissertation is that when there are many credences which are reasonable, probability sentences are vague. And the best approach for interpreting vague sentences is the technique of supervaluations. The vagueness in probability sentences mirrors the permissible imprecision in reasonable degrees of belief.

In this chapter I'll defend the use of supervaluational semantics for probability sentences against a number of recent attacks. The only change I make to the orthodoxy here is that I claim that on a supervaluational account, sentences containing vague terms have scope ambiguities. Orthodox accounts confuse conventional resolutions of this vagueness with the literal or semantic content of sentences. Even if my general defence against these attacks fails, there are grounds for thinking the supervaluational approach is more plausible when it is applied here than when it is applied generally to vagueness.

In section 2 I'll set out an orthodox supervaluational account. The orthodoxy may be (slightly) wrong, but it is clearly the best way to explain the mechanics of a supervaluational account. In section 3 I'll show how this applies to probability sentences. Williamson objects to

supervaluational accounts because, *inter alia*, they don't respect Tarski's T-schema. I think this is a sound criticism of the orthodoxy; and in section 4 I introduce an amendment to overcome it. This amendment allows for sentences which would ordinarily be supervaluated – e.g. sentences containing vague terms, sentences containing definite descriptions – to have a scope ambiguity. Section 5 discusses the notion of validity in the context of sentences without truth values. Sections 6 and 7 respond to other recent objections to supervaluational accounts. Sections 8 and 9 discuss the notion of reasonable degrees of belief that I have used in the earlier sections. Section 8 fleshes out the arguments discussed in chapter 3 for there being more than one reasonable probability function. And in section 9 I argue that it is reasonable (though not epistemically mandatory) to have precise degrees of belief in all propositions.

5.2 Supervaluations

The simplest introduction to supervaluational semantics is through its account of vague predicates. Assume Jack is such that it's unclear whether he is tall (for an adult male)¹ or not. Jack, for example, might be somewhat taller than the average adult male, but not a lot taller. Neither 'Jack is tall' nor 'Jack is not tall' would seem like appropriate things to say. We might just say it's vague whether or not Jack is tall.

Many people have believed we should say that 'Jack is tall' lacks a truth value. The supervaluationist has a principled reason for saying this, as well as a full story about how 'Jack is tall' behaves in logical compounds. The core of the story is the set of possible precisifications of *tall*. More generally a precisification will make all of language precise, but we'll leave that aside for now. A precisification V of tall assigns truth values (we're assuming here that *true* and *false*

¹ Whether or not someone is tall of course depends on the context. Someone can easily be tall for a jockey and not tall for a basketballer. We'll assume context is invariant in what follows.

are the only truth values) to every sentence of the form ‘ x is tall’ in accordance with three constraints.

First, all (first-order)² truths of English must be preserved in the precisification. So if Al really is tall, then on every precisification ‘Al is tall’ must come out true. Similarly, if Bill really is not tall then on every precisification ‘Bill is tall’ must come out false. Secondly, some conceptual truths must be preserved. These are what Kit Fine (1975) calls penumbral connections. If Jack is taller than John, there can’t be a precisification according to which ‘Jack is tall’ comes out false, and ‘John is tall’ comes out true. Even if both Jack and John are borderline cases of tallness, so there are precisifications on which ‘Jack is tall’ comes out false and precisifications according to which ‘John is tall’ comes out true, there can’t be a precisification on which both occur, because no tall person is shorter than a not tall person. Thirdly, truth values of compounds in the precisification are given by the usual rules. So ‘Jack is not tall’ comes out true iff ‘Jack is tall’ comes out false; ‘Jack is tall and John is true’ comes out true iff ‘Jack is tall’ comes out true and ‘John is tall’ comes out true, and so on for the other connectives.

There will be uncountably many precisifications, but they will have some things in common. As we noted ‘Al is tall’ will come out true on all of them if Al really is tall. ‘Jack is tall and Jack is not tall’ will come out false on all precisifications, even if the conjunct which is false differs for different precisifications. If Jack is taller than John, ‘John is tall and Jack is not tall’ will also come out false on all precisifications, with again the possibility that different conjuncts come out false on different precisifications. The supervaluationist defines truth (in English) to be truth on

² I am using ‘first-order’ in the sense developed by Tarski. So ‘Snow is white’ is a first-order sentence, ‘“Snow is white” is true’ is a second-order sentence, and so on. Section 6 will discuss the need for this restriction.

all precisifications, and falsity (in English) to be falsehood on all precisifications. These will be referred to as supertruth and superfalsity.

It is worth pausing to note some of the effects of this method of evaluating compounds. There can be conjunctions which are false despite having no false conjuncts. Conversely, there can be disjuncts which are true despite having no true disjuncts. This will be the case when for every precisification, one or other disjunct is true, but there is no disjunct true on all precisifications. To take the classic example, ‘Jack is tall or not tall’ comes out true even though neither ‘Jack is tall’ nor ‘Jack is not tall’ come out true.

A similar effect arises for existentially quantified sentences. There can be predicates F , G such that ‘Some F is a G ’ is true even though there is no F , say a , such that ‘ Fa and Ga ’ is true. This will occur when on every precisification some F or other is G , but there is no F which is a G according to all precisifications. Conversely, ‘All F are G ’ can be false even if there is no a such that ‘ Fa and Ga ’ is false. These general points will be important in what follows, but as I’ll show in section 4 there is need for great care in how they are interpreted.

5.3 Probability Sentences

As discussed in chapter 4, some probability functions are reasonable and some are not. There are still some matters to discuss concerning the distinction, but I’ll assume for now we have it. Each of these functions plays the same role in my theory that precisifications play in supervaluational accounts of vagueness. The truth value of a sentence according to each function can be easily worked out. A sentence like ‘The probability of p given q is greater than $1/2$ ’ is true according to a probability function Pr iff $Pr(p \mid q) > 1/2$. Probability sentences are true *simpliciter* iff they are true on all precisifications, false iff they are false on all precisifications.

This approach has much pragmatically to recommend it. First, all the theorems of the classical probability calculus turn out to be true. So, for example, it is true that ‘For all p , q there is an x

$\in [0, 1]$ such that the probability of p given q is x '. Given the great amount of work that has been put in over the years into developing the probability calculus, it would be unfortunate to have to give it all up, or to say that it is inapplicable to actual uses of 'probability'. The objections to unwarranted precision in earlier chapters do not show that we should be rid of such theories; rather they show we should be rid of careless interpretations of them.

In particular it is not the case that there need be an x which makes 'The probability of p given q is x ' true. If we'd ordinarily say the probability of p given q is vague, say vague over the interval $[0.4, 0.55]$, then for some x the sentence 'The probability of p given q is x ' will be neither true nor false (i.e. for those x in $[0.4, 0.55]$) and for all other values of x it will be false. However, the existentially quantified sentence will be true because on all precisifications (all reasonable probability functions), the probability of p given q takes a precise value.

Some might view this preservation as a negative: it isn't just an interpretation of some of the theorems of the probability calculus they wanted to throw out. There are, however, other indisputable advantages. Say the probability of p (I'll omit references to evidence except where necessary) is vague over a large interval, and the probability of r given p is some high value, say 0.9. Hence, the intervals over which the probability of p and the probability of p & r are vague overlap. So, for example, the probability of p might be vague over $[0.4, 0.55]$, and the probability of p & r might be vague over $[0.36, 0.495]$. The probability of p is clearly greater than that of p & r . On a supervaluationist account this is true, since according to every reasonable probability function we have $Pr(p) > Pr(p \ \& \ r)$. On a purely 'intervalist' account I doubt this can be shown, for reasons given in 3.5.3.

So we have three reasons for adopting the supervaluational semantics suggested here for probability sentences. First, it allows us to keep the theorems of the classical probability calculus. Secondly, it explains why we might have thought some of these theorems were inappropriate by

showing that some natural interpretations of these theorems are false. Thirdly, it gives the intuitively appropriate truth value to comparatives, something that other semantics which allow vagueness cannot do. These benefits are distinct from the general benefits of adopting supervaluational semantics for definite descriptions or vague terms, such as the natural account it provides of compounds.

There is one other complication which can be resolved by a supervaluational approach. I have said in the above that all probability sentences make a reference, either explicit or implicit, to an evidence set. When this is implicit there might be some vagueness as to what the evidence is. I think the best solution here is again a supervaluational approach. So a probability sentence is true iff it is true on all possible resolutions of this vagueness, false iff it is false on all possible resolutions of this vagueness, and lacking a truth value otherwise.

When talking about the probability of a proposition about the future this is not normally a problem, as usually this is a reference to chance. So, if I say that the probability of Al Gore winning the 2000 election is x , I mean the chance of him winning is x , where chance is defined as in the previous chapter. However, for propositions about the past this cannot generally be right, because the chance of the proposition being true will be zero or one. So if I say ‘The probability the suspect is guilty is at least $1/2$ ’ it mightn’t be clear whether the implicit evidence set is all the evidence in the public domain, all the evidence I have, all the evidence that I share with my hearer, or some other set. Context will usually narrow the range a little, but it can’t be guaranteed to do all the work. In these cases what I say might be true on some resolutions of the vagueness and false on others, hence lacking a truth value.

For these reasons I agree to some extent with the claim in Price (1984) that probability sentences can fail to be truth-apt (which I interpret as lacking a truth value) because the reference to evidence may not be precise. He says that probability sentences will be truth-apt when there is agreement

as to what the evidence is³, but when there is no such agreement there can be disputes where speakers say opposing things without either necessarily being mistaken, so their claims are not truth-apt. The position adopted here is that when the different implicit evidence sets speakers use are close enough without being exactly the same, there can be genuine disputes, and hence truth-apt claims. For example, I can have a debate with someone about whether Oswald probably killed JFK even if we disagree a bit about what counts as evidence, so we can at least presuppose that our utterances are truth-apt. What counts as ‘close enough’ in the above definition will depend crucially on what is being talked about; for uncontroversial claims it will be defined widely and, conversely, for borderline claims it might be defined narrowly.

One advantage of this position is that we do naturally assign truth values to some probability sentences even when the implicit evidence is a bit imprecisely defined. For example, if I say ‘There is probably no present king of France’, even if it is massively unclear from context what is to count as evidence, the utterance seems to be true. On the other hand for more disputable claims, such as in the debate about Kennedy’s assassin, imprecision in the definition of evidence may destroy the truth-aptness of the claims. All these results sound intuitively plausible, at least to my ear.

5.4 Scope and the T-schema

Timothy Williamson (1994) claims adopting that supervaluational semantics leads to a denial of Tarski’s T-schema. And if the T-schema isn’t all we know about truth, it is, claims Williamson,

³ I disagree with this in general, because of the possibility of imprecision, but since that isn’t what’s at issue just here, I’ll assume all probabilities are precise. I also presume that agreement as to what the evidence is need only be with an imaginary interlocutor, so that a probability sentence will be truth-apt if we could specify what the evidence is, even if this doesn’t happen in the actual conversation. Should this not be the case then the truth value of speakers’ utterances will depend implausibly on subsequent utterances of actual interlocutors.

a large enough part that we shouldn't give it up. The point of this section is to make two responses to Williamson, and to show that at least one of these allows a response to an objection which Mark Sainsbury believes refutes supervaluational approaches.

Say F is a vague predicate. On a supervaluational approach, the sentence ' $Fa \vee \neg Fa$ ' will always be true, for on every precisification of F one or other disjunct is true. However, the sentence "' Fa ' is true or "' $\neg Fa$ ' is true', need not be true. In particular, if a is a borderline case of being an F , then on some precisifications Fa will be false on the precisification, so Fa is not supertrue. Further, there will be precisifications according to which Fa is true on the precisification, so $\neg Fa$ is not supertrue. So we can have ' $Fa \vee \neg Fa$ ' true without "' Fa ' is true or "' $\neg Fa$ ' is true' being true. This result is acknowledged by supervaluationists; indeed something like it is usually taken to be one of the distinctive claims supervaluationists make.

Now if this is right, Williamson notes, we can't keep the T-schema. For that schema says that for any proposition A , the material biconditional ' A iff " A " is true' holds. Now, given this, and some fairly straightforward logical machinery⁴, we can deduce the material biconditional (1), where A and B are any propositions:

$$(1) \quad (A \text{ or } B) \text{ iff } ("A" \text{ is true or } "B" \text{ is true}).$$

Strictly, (1) isn't an instance of the T-schema, but it is an implication of it. As noted above, if the T-schema is taken to be incontrovertible, the supervaluationist has a difficulty. It is a difficulty Williamson thinks is inescapable without moving to an epistemic conception of vagueness, according to which there is only ever one permissible precisification, but we are ignorant as to

⁴ To be precise, the standard natural deduction rules \vee -introduction, \vee -elimination and *modus ponens*.

what it is. I will make two responses; the first of which is very similar to one Williamson discusses and rejects, and the second of which is quite different.

The response which Williamson discusses is to claim that there is a disquotational truth predicate, $true_T$. This response was first raised in Fine (1975), though he develops it somewhat differently to the way I am doing. “‘ A ’ is true’ is $true_T$ according to a precisification iff A is true according to that precisification, where *iff* is read as the material biconditional. Williamson acknowledges this move rebuts his objection, but claims that once it is adopted there is no reason not to identify truth with $truth_T$ rather than with supertruth. For example, a plausible truth-functional account of validity can be given if this identification is made⁵. And if we do identify truth with $truth_T$ rather than with supertruth, Williamson claims we lose all that is distinctive about the supervaluationist position. We just fall back into his preferred position, that vagueness is an epistemic phenomenon. “Of supervaluationism, nothing remains articulate,” he concludes (1994: 164).

A different way of running this response is to claim that there is no reason *supertrue* shouldn’t behave like $true_T$. The standard way of assessing sentences like “‘ A ’ is true or ‘ B ’ is true’ is to evaluate each disjunct separately. But this goes against a central tenet of supervaluationist theory, which says that to determine the truth value of complex sentences containing vague terms we

⁵ Williamson notes that the supervaluationist might argue that identifying truth with $truth_T$ will mean that not every sentence will be definitely true or definitely not true, and we might have thought this a flaw in a semantic theory. However, he claims that given the existence of higher-order vagueness, even identifying truth with supertruth will not preserve this condition. And here some card-carrying supervaluationists agree with him. See for example Williams (1976).

For what it’s worth, we can also give a plausible truth-functional account of validity if we identify truth with supertruth, but validity won’t just mean truth preservation; it must also mean never going from non-false premises to false conclusions. More on this in later sections.

have to look at the truth value of the whole sentence according to each precisification. This rule, it seems, needs to be extended to sentences which contain vague terms *quoted*.

On this variant of supervaluationism, what ought we say about the truth of “‘*A*’ is true’ according to a precisification? According to that precisification, all terms in *A* are precise, and hence all have just one admissible precisification. Hence this sentence will be true (according to all precisifications) according to that precisification iff it is true according to that precisification. In a more digestible soundbite, supertruth according to a precisification equals truth according to that precisification. Hence (1) will hold. In the troubling case, it might be impossible to identify the true disjunct on the right of the biconditional, but there will be one.

If this move is made, it becomes impossible to *say* what is distinctive about the supervaluationist approach. To see this, consider a case where we would normally say that *Some Fs are Gs* is supertrue while for any particular *F*, say *x*, *Gx* would not be supertrue, as different *Fs* are *G* according to different precisifications. All of the following are literally false according to this approach.

- (2) There is no *x* such that ‘*Fx* and *Gx*’ is supertrue.
- (3) There is no correct answer to the question, ‘Which *F* is *G*?’.
- (4) No *F* is such that it is a *G* according to all precisifications.

All of these have the common feature that they are true if we are speaking from outside the supervaluational framework; speaking about that framework without using it. But this cannot be done. So says the hard-line supervaluationist. The attempt to say things like (2) to (4) and hence capture what is distinctive about the supervaluational position is an attempt to find a linguistic ‘view from nowhere’. Williamson is right to say that nothing articulate and distinctive remains about this supervaluational position.

It isn't clear, however, that there remains nothing distinctive at all about this approach. It can't be said what it is, but perhaps this isn't too surprising. We shouldn't be too surprised to learn that there are things about a language that we can't say in that language. That will hold whether the language is as precise as the formal language of *Principia Mathematica*, or as vague as ordinary English. We might hint at what is distinctive about the supervaluational position by saying things like (2) to (4), but if this variant of supervaluationism is correct, vague languages are rather incomplete in the sense that we cannot use them to say things we would like to be able to say. The idea is that by saying (2), (3) or (4) we are exploiting the circumlocutional nature of our utterances to implicate (i.e. point to) something which can't be said⁶.

I think this position has some attractions (theoretical simplicity for example) but I can't imagine it will persuade anyone. There is a different way. It might be permissible to read (2) to (4) in a way in which they come out literally false, but there is also a reading under which they are true. All of these sentences have scope ambiguities, and different resolutions of these scope ambiguities will lead to differing truth values for the sentences. (2), for example, could be rendered as (2a) or (2b), the first false, the second true.

- (2a) According to all precisifications, there is no x such that ' Fx and Gx ' is true (according to all precisifications).
- (2b) There is no x such that according to all precisifications ' Fx and Gx ' is true (according to all precisifications).

⁶ Interpreted literally, there is a problem with this last paragraph. On every precisification there is nothing that can't be said, so the sentence 'There is something that can't be said' must come out false by virtue of being false on all precisifications. I take this to be another confirming instance of Ramsey's rule: what can't be said can't be whistled.

Once it is seen that there is always going to be a reading like (2a), where the ‘internal’ reading⁷ of the sentence is formed by prefixing ‘according to all precisifications’, we see that these scope ambiguities are ineliminable. So there is no point trying to eliminate them by adding more and more words. Rather we use a convention such that the quantification over precisifications is given wide scope in ‘simple’ sentences and narrow scope in ‘complex’ sentences. So according to the convention, (2) should be read as (2b)⁸.

The scope ambiguity isn’t confined to sentences like (2). On the contrary, ‘Some *F*s are *G*s’ is capable of being read both like (2a) and like (2b). Here, though, the convention is that we read it like (2a). Now every instance of the T-schema will hold as long as we disambiguate both sides in the same way. The argument purporting to show that, from premises the supervaluationist is committed to, an instance of the T-schema is falsified rests, according to this analysis, on a fallacy of ambiguity.

Sainsbury (1988: 41) claims that supervaluationists are committed to denying the existence of vague terms. Consider the classic vague term *heap*. Every precisification of *heap* can be

⁷ In contrast to the ‘view from nowhere’ reading.

⁸ This convention could be read as a conventional implicature. This would lead to the conclusion that the T-schema was literally true, as the left and right sides have the same semantic content, but they have differing conventional content. In any case, it is uncontroversial that the different sides have different pragmatic properties.

There might be a question as to how such a convention could have come about. My best explanation would be that because the developers of supervaluational semantics said things like (2), (3) and (4) and intended them to be read truly, readers applied the principle of charity and read them as truths. When subsequent writers talking about supervaluational semantics use sentences like these, we are reminded of the classical usage, and hence interpret them the same way. Thus a convention develops, and it spreads again by usage of charity on textbook writings using similar sentences in a similar way. I make no claim for the empirical accuracy of this story: however, I do hope that some story like it can be told to explain the creation of the convention.

summarised by a single number n : piles of sand with n or more grains are heaps, and with less than n grains are not heaps. So Sainsbury notes according to supervaluationist accounts, (5) should be true.

(5) For some number n , an n -grained collection is a heap, but a $(n - 1)$ -grained collection is not.

This, he claims, is pretty much a denial of the vagueness of *heap*; vagueness means *inter alia*, having no sharp boundary. The scope ambiguity analysis allows the supervaluationist a reply. (5) is ambiguous between (5a) and (5b); (5a) is the reading which is literally true and (5b) is the reading which amounts to a denial of the vagueness of *heap*. But no reading is true and amounts to denial of the vagueness of *heap*.

(5a) According to all precisifications, for some number n , an n -grained collection is a heap, but a $(n - 1)$ -grained collection is not.

(5b) For some number n , according to all precisifications, an n -grained collection is a heap, but a $(n - 1)$ -grained collection is not.

So the supervaluationist has two replies to the arguments put by Sainsbury and Williamson. The first involves the rather startling claim that once we supervaluate many of the analyses of vague sentences, we find they are literally false. There are, on this reply, many things we would like to be able to say about sentences with vague terms, but our language doesn't have the resources to say them. The second reply is that sentences containing existential quantifications or disjunctions involving vague terms have scope ambiguities. The conclusions to which Williamson and Sainsbury believe the supervaluationist is committed can only be drawn by failing to distinguish carefully enough the different readings of these sentences.

There is an important objection to this view which needs a rather careful argument to refute. I have held that disjunctive utterances containing vague terms are ambiguous between their two supervaluational readings, and this ambiguity is a scope ambiguity. There is no reason why my claim shouldn't extend to ordinary utterances. So, assuming again that Jack is a borderline case of baldness, then (6) should be ambiguous between a reading on which it lacks a truth value and a reading on which it is false.

(6) Jack is bald.

Now there's nothing too implausible about the first part of this conclusion. *Ex hypothesi*, Jack isn't one of the bald things, so we should be able to say (6) is false. On the other hand, since Jack is a borderline case of baldness, many writers have thought it correct to deny (6) a truth value. So far, then, we don't have a telling objection. What is a problem is that it looks implausible to say that the ambiguity in (6) is a scope ambiguity. A scope ambiguity requires that there is a quantifier which could be placed in various parts of the sentence; (6) doesn't have the required multiple insertion points and, hence, is not susceptible to scope ambiguity. I could avoid this problem by postulating a brute ambiguity; but this is incredible and in any case would be inconsistent with the theoretical position I adopted in the last chapter. Apart from homonyms like 'bank', all ambiguity should be explainable in principle; and when a word behaves ambiguously in several contexts, the same explanation should be given for the range of behaviour displayed.

To show that the ambiguity in (6) really is a scope ambiguity we need to consider why a sentence lacks a truth value. To consider this we must first consider why the sentence might be true. As Dummett (1959) stressed, it is a mistake to consider truth in isolation from our practices of asserting and believing. We like to believe true sentences; generally we like to assert true sentences. We want to make arguments which preserve truth. Any purported account of truth, belief, assertion and arguing which does not take these things into account will be mistaken. If

we cash out these desires as beliefs about what would happen in ideal conditions (truth is what we believe in ideal conditions) we are lead into constructivism, but there is no need to analyse them that way. Lewis (1988) has showed there are theoretical reasons for not analysing desires as beliefs of any kind. We can accept these points about truth without becoming anti-realists.

I have given some properties of truth; we have certain desires expressible in terms of it and it has a certain place in a theoretical network that includes belief, assertion and argument. Assume I can give enough properties to get an analysis of truth. So a sentence is true iff it has property *T*. The analysis is a little misleading; the sentence is *T* because it is true, not *vice versa*. It must now be decided what will count as an analysis of ‘false’. If I was committed to saying all sentences are true or false I could simply say that a sentence is false if it is not *T*. Given that I am trying to analyse what is meant by saying a sentence lacks a truth value, this looks like the wrong path to take.

Rather, I will say that a sentence is false iff its negation is *T*. This seems to capture the motivation of many theorists who deny some sentences truth values; they fear that saying a sentence is false commits them implausibly to the truth of some other sentence. Since that other sentence is not one they want to assert or believe, not one they would like to have as the conclusion of a good argument with true premises, they deny the original sentence is false.

This raises the rather obvious question, how do we tell what the negation of a sentence really is. In the case of (6) this is somewhat non-trivial. Many would argue that the negation of (6) should just be (7).

(6) Jack is bald.

(7) Jack is not bald.

Since (7) is not *T*, and it is the negation of (6), we conclude (6) is not false. The problem is that the second premise here looks dubious. If we accept a Russellian account of definite descriptions, we have to say that (9) is not the negation of (8).

(8) The present king of France is bald.

(9) The present king of France is not bald.

Famously, Russell said they were both false⁹. Yet they stand in the same syntactic relationship to each other as (6) and (7), so if negation is a syntactic relationship, (7) is not the negation of (6). I don't take this little argument to show that (7) is not the negation of (6); I do take it to show that the relationship between (7) and (6) isn't as clear as some have supposed¹⁰. The official definition of the negation of (6) in introductory logic texts is generally (10).

(10) It is not the case that Jack is bald.

(11) 'Jack is bald' is not true.

I take it that it is unclear whether (10) should be interpreted as (7) or as (11). The good news from the view of the ambiguity theory being defended here is that (10) does seem to have a scope ambiguity, and the two readings correspond to (7) and (11).

(7a) According to all precisifications it is not the case that Jack is bald.

(11a) It is not the case that Jack is bald according to all precisifications.

This explains why (6) is ambiguous. The truth value of a sentence depends both on whether it is *T* and on whether its negation is *T*. Since (6) is not *T*, it cannot be true. But whether it is false or

⁹ At least on one reading of (9).

¹⁰ Including Russell himself. He said that in cases of vagueness, where we deny 'Jack is bald' and deny 'Jack is not bald', we are committed to denying a case of excluded middle (1923: 88).

gappy depends on whether its negation is *T*. There is a scope ambiguity in its negation, so we can say that the truth value of (6) is ambiguous. Finally, whenever the truth value of a sentence is ambiguous, that sentence is ambiguous. Hence the ambiguity in (6) is not ‘brute’, it can be explained in terms of scope, so the objection fails.

This analysis of negation is not at all new; something like it appears to be going on in Russell (1905: 53). He says that (9) is ambiguous; whether or not it is true depends on whether we regard the denoting expression as, in his terms, ‘primary’ or ‘secondary’. It is false if it says the present king of France is among the things which are not bald; this is the primary reading. It is true if it denies that among the bald things is the present king of France; this is the secondary reading. The main difference is that it is clear that the negation of (8) is the secondary reading of (9), whereas in this example it does look genuinely ambiguous what should be the negation of (6).

For ease of reference, I’ll call the first of the two responses listed here *strong supervaluationism*, and the second *moderate supervaluationism*. The traditional approach is then called *weak supervaluationism*. The names derive from how far the various theories say supervaluational approaches should be applied. Since the strong approach seems so unlikely to be persuasive, unless otherwise stated I’ll assume the moderate version is correct in what follows.

5.5 Validity

Williamson’s other primary objection to the supervaluationist approach to the resolution of vagueness is that it lacks a plausible account of validity. Within the supervaluational literature there are two definitions of validity: local and global validity. Local validity says that an argument is valid iff it is traditionally valid on every precisification. In other words, if whenever all the premises are true on a precisification, the conclusion is also true. Global validity says that an argument is valid if whenever the premises are all true on all precisifications, the conclusion is

also true on all precisifications. In other words, validity is supertruth preservation. Local validity is defended by Dummett (1975), global validity by Fine (1975).

Williamson claims that since the supervaluationist identifies truth with supertruth, they must plump for global validity as the appropriate account of validity. He then notes that this account of validity will not allow some intuitively plausible inference rules (in particular *reductio ad absurdum*) and concludes the supervaluationist has some implausible commitments.

I'll concede that global validity is an implausible standard of validity. However, I don't see why the supervaluationist should be committed to it. Sentences of the form ' p entails q ' (or alternatively ' Γ entails q ') have a scope ambiguity of the sort mentioned above. We can interpret (1) as either (1a) or (1b).

- (1) p entails q .
- (1a) According to all precisifications, p entails q .
- (1b) p according to all precisifications entails q according to all precisifications.

Moderate supervaluationism says that (1a) is at least a legitimate reading of (1). Indeed, since it is the only reading licensed by strong supervaluationism, it is perhaps the primary reading. If we insist on reading (1) as (1b) then we may be lead to incoherence; but if we insist on supervaluating some sentences and not others this is perhaps not too surprising.

If we interpret (1) as (1a) we get Dummett's local validity as the proper account of validity for supervaluationists. Since local validity preserves all the standard rules of entailment Williamson's argument would collapse. We don't, when using local validity, use our definition of truth as supertruth in our account of validity. Williamson suggests that this implies we are not really committed to this definition. Were we to read (1) as (1b) and still use local validity as the

appropriate criteria Williamson would be correct here, but, as that's not what we're doing, I don't see the objection here.

Even if we interpret (1) as (1b), global validity is still not the right theory of validity. Such an approach assumes that, in effect, we have three truth values¹¹: <true, gap, false> and only the first of these is designated. Such an approach runs foul of an objection Dummett (1959) makes to Strawson's account of definite descriptions. Strawson (1950) says that 'The present king of France is bald' lacks a truth value because it has a non-referring definite description. Dummett replies that Strawson hasn't accurately distinguished lacking a truth value from being false. In particular, since 'gap' is in effect undesignated in Strawson's system, it looks like Strawson has merely distinguished two types of falsehood rather than distinguished falsehoods from sentences lacking truth values.

To motivate a theory which says that some sentences lack truth values, I have to do something which distinguishes 'gap' from 'false'. I made some progress towards that in the previous section; false sentences have true negations, gappy sentences do not. I can now make a more direct distinction between the two¹². 'Gap' is, in the best account of validity for three-valued logics, neither a designated nor an undesignated value, as the account of validity is not a designated value account. Rather, I say that an argument is valid iff it satisfies the following criteria.

' p entails q ' is valid iff

- (i) it is impossible for p to be true and q to be not-true; and

¹¹ Strictly 'gap' is not a truth value, it is the lack of one, but spelling this out every time would make the exposition unreadable, so I presume the reader can remember this caveat.

¹² Neither distinction I am using is redundant. The previous section showed how to determine whether particular sentences are gappy or false; this section provides a stronger argument for the claim that there is a real difference between these two 'values'.

- (ii) it is impossible for p to be not-false and q to be false.

It is easy to check this cannot be represented as a designated values account of validity. Such an account is quite popular in the literature¹³, and it's clear why. Not only does it escape Dummett's objection to Strawson, it preserves contraposition and *reductio*. ' p entails q ' is valid iff ' $\neg q$ entails $\neg p$ ' is valid. And ' p entails q and $\neg q$ ' is valid, iff for any r , ' r entails $\neg p$ ' is valid. The motivation for this account is the scale of truth values <true, gap, false> mentioned above. An argument is valid iff it is impossible in moving from the premise to the conclusion to move 'down' the scale, either from true to anything else, or from gap to false.

There are some complications in the story for when we have a set of premises or a set of conclusions. And this story will not preserve all the classical rules. For example conditional proof is no longer sound; nor is argument by cases, sometimes called \vee elimination. However, the fact that even on my non-preferred reading of (1) I can keep *reductio* and contraposition suggests that there is no objection from this direction to the supervaluational approach.

5.6 Models and Conceptual Truths

Fodor and Lepore (1996)¹⁴ have recently claimed there is a simple objection to supervaluational semantics. They claim that the precisifications which are the core of this account do not satisfy a mandatory condition on models for a language: that all conceptual truths are true in the model. Because of this there is no reason to think 'truth in all precisifications' is really truth.

The objection is built up as follows. Vague terms like 'bald' have a penumbra. Assume for the sake of the argument that a person with 1/9 of their head covered with their own hair is in this penumbra, neither definitely bald nor definitely not bald. If the ratio is wrong, it can be changed.

¹³ Most recently McDermott (1996); see the references therein for earlier proponents.

¹⁴ All references in this section to this paper unless otherwise stated.

All that matters is that there is one, which everyone agrees. Let S be the sentence ‘A person with a head-to-hair ratio of $1/9$ is bald’. If our assumption is right, S is neither definitely true nor definitely false. Indeed, on standard accounts S is neither true nor false *simpliciter*.

That S is not definitely true is not something that we discovered by looking at the world. Of course we discovered what S means by looking at the world, but once we found that out we worked out by conceptual analysis alone that it is not definitely true. So ‘ S is not definitely true’ is a conceptual truth. Fodor and Lepore then wield what they call principle (P).

“(P) Conceptual truths must be respected by all classical models, including classical valuations.” (521).

The justification is that a purported model of a language which does not respect conceptual truths is not a genuine model. “If there are conceptual truths, then they determine what the topic under discussion *is*, so they must not be flouted on pain of equivocation.” (521). But precisifications do not satisfy this criteria. To see this, note that in many precisifications, ‘ S is not definitely true’ comes out as false, despite being a conceptual truth. In all those precisifications in which it does come out as true, the conceptual truth ‘ S is not definitely false’ comes out as false. So there are no precisifications which satisfy principle (P).

As well as this general argument for (P), they have an *ad hominem* against the supervaluationist who does not accept it. Say, Al and Bill each have $1/9$ of their head covered with their own hair. Then there are, according to supervaluationism, acceptable precisifications according to which ‘Al is bald’ comes out true. There are also acceptable precisifications according to which ‘Bill is bald’ comes out false. But there is no acceptable precisification according to which ‘Al is bald’ comes out true and ‘Bill is bald’ comes out false. Referring to Fine (1975), Fodor and Lepore claim that the supervaluationists’ reason for this is that precisifications must preserve conceptual truths, in this case the conceptual truth that baldness supervenes on head-to-hair ratio. So by their

own lights, supervaluationists are committed to (P). But there are no precisifications which are acceptable according to (P).

The response to Fodor and Lepore is three-fold. First, some reasons for thinking that (P) need not be satisfied by models are discussed. The basic point is that precisifications of English are not meant to be meaning preserving at the level they are discussing. It's no news to say 'bald' in a precisification means something different from 'bald' in English because the former is precise and the latter is vague. The second is that we can distinguish acceptable from unacceptable precisifications without relying on (P). Finally, it is argued that there is no way to make sense of Fodor and Lepore's positive suggestion, which is that *S* is gappy, and must be so on all models, without using supervaluations. In sum, I don't dispute that giving up principles like (P) is part of the cost of adopting a supervaluational account, but I think the cost can be shown to be rather small, and the benefits rather large.

Fodor and Lepore treat precisifications as languages, so we can talk about the meaning of a word in a precisification. It is simpler to treat them as sets of true sentences, or equivalently as (complete) functions from sentences to truth-values. In any case, we ought to be able to determine the meaning of a word in a precisification from the set of sentences containing that word which are true. Let E be the set of first-order truths of English. So 'Snow is white' is in E , but "'Snow is white' is true' is not in E . Let E^* be any maximally consistent superset of E which is closed in the following ways:

$$A \ \& \ B \in E^* \text{ iff } A \in E^* \text{ and } B \in E^*$$

$$A \ \vee \ B \in E^* \text{ iff } A \in E^* \text{ or } B \in E^*$$

$$\neg A \in E^* \text{ iff } A \notin E^*$$

$$\text{If } A \rightarrow B \in E^* \text{ and } A \in E^* \text{ then } B \in E^*$$

$$'A' \text{ is true} \in E^* \text{ iff } A \in E^*$$

E^* is a precisification of English iff it satisfies all of these conditions. I'm intending ' \rightarrow ' here to be read as a natural language conditional; the condition regarding it is redundant if it is read as a material implication. This definition is intended to perform two jobs. First, any (first-order) truths of English are true in all precisifications. So if Jack is bald then 'Jack is bald' will be an element of all precisifications.

Secondly, precisifications preserve what Fine called the penumbral connections, like 'Taking someone's hair away doesn't change them from being bald to not-bald'. The way these were preserved by Fine suggested that supervenientists were committed to (P). However, here they are preserved not because they are conceptual truths, but because they are first-order. For example, 'Baldness supervenes on hair-to-head ratio' is a first-order truth, so it will be in E^* . Returning to our example of Al and Bill, this implies that 'If Al is bald, Bill is bald' is in E and hence E^* . I'm assuming here that if A entails B then 'If A , B ' is true. Hence there can be no precisification in which Al is bald and Bill is not bald. Similarly, we can find general (perhaps conceptual) first-order truths which imply that bald people can't have a higher hair-to-head ratio than non-bald people.

It's a trivial fact that for one object to model another, it doesn't have to have all the properties of the object being modelled, or indeed all the essential properties. Consider the use of crash test dummies to model the behaviour of humans in car crashes. So on a natural reading of 'model', there is no reason to say that precisifications are not models of English just because they lack essential properties of English. There are two ways in which breach of (P) by precisifications would cause problems, but neither seem to be realistic possibilities.

First, if someone were claiming 'bald' in E^* means the same as 'bald' in English, then breaches of (P) would be problematic. Meaning-preserving translations ought to preserve conceptual truths. But there is a bigger problem with this approach. It would imply that we can work out the

meaning of ‘bald’ by just stipulating a cut-off point. Since any stipulation would provide the meaning, this would lead to blatant inconsistencies. This clearly isn’t what supervenientists are trying to do. The meaning of ‘bald’ isn’t given by its behaviour in a particular precisification, but in the set of them.

Secondly, there might be a difficulty if there were permutation problems. Granted that E^* is a model (not an analysis or translation) of English, we have to determine which English words are being modelled by particular words in E^* . Ideally there will be a function from words in E^* to words in English. However, there might be multiple plausible functions. Were this to occur then E^* wouldn’t be a good model, and there might be wholesale difficulties for the supervenientist, because it wouldn’t be clear if the equivalent sentence in the model to a particular sentence of English were true or not. However, there is little evidence that there are such problems, and hopefully the preservation of all first-order truths in all precisifications prevents such a difficulty occurring. If there is a problem on these lines, no objector has yet shown it.

From this we can determine that the status of ‘Jack is not definitely bald’ in a precisification will depend on how we read ‘definitely’. If we read it as a function from predicates to predicates (like ‘very’) we will assume that this is a first-order truth (if Jack is in the penumbra of bald), and so it will be true in all precisifications. On the other hand, if by this we mean “‘Jack is bald’ is not definitely true’, then it is a second order truth. This will be true in some models and not in others. I assume that what is true in a model is, for reasons given in previous sections, definitely true in that model. So this conceptual truth will not be preserved in all models, but we have reasons for thinking models need not preserve conceptual truths.

This answers the letter of Fodor and Lepore’s objections. Precisifications need not obey (P) because they are models of English, not analyses or translations of it. And we can capture Fine’s penumbral connections without relying on (P), so there is no *ad hominem* argument for (P). There

remains, however, a powerful related problem. Why, given that precisifications are merely models for English, should truth on all models mean truth in English? For the first-order sentences, the answer is trivial. If a first-order sentence is true on all precisifications which preserve all first-order truths, then of course it is true. The question is why this should be the case for the higher-order sentences.

There is here no quick answer. The long answer turns on four points. First, given that this approach works for first-order sentences, there is an argument from theoretical simplicity to use it everywhere. This needs little elaboration, but it isn't that powerful on its own. The second is that only supervaluational theories justify saying that some sentences lack a truth-value. The third is that, given that some sentences lack a truth-value, only the moderate supervaluational approach avoids implausible results. A fourth, which won't be elaborated but which might have some power, is that the supervaluational explanation of the truth-value of compounds with gappy components is preferable to its rivals.

Many of these points have been stressed in previous sections, and there's no need to recap them here. I have made much of the general argument against the existence of 'gappy' sentences in my replies to some objections of Williamson. We can't just stipulate that sentences without truth values exist, we need to explain what we mean by this, and in particular we need to explain why we take being 'gappy' to be different from being 'false'.

Even if this general hurdle is cleared, we need to explain why some vague sentences should be 'gappy'. I have argued above that the easiest way to do this is to show that the negations of these sentences are not true. If it can be shown that 'Jack is not bald' should be considered the negation of 'Jack is bald', the task might be complete. The supervaluationist can show this easily; 'Jack is not bald' is undoubtedly the negation of 'Jack is bald' on all precisifications, so it is plausibly its negation *simpliciter*. For other theories, however, the arguments near the end of section 5.4 seem

to show this can't be done. So Fodor and Lepore's preferred position, that sentences like 'Jack is bald' are gappy and that's all there is to the matter, looks flawed.

Moderate supervaluationists can explain how there can be gappy sentences even though Tarski's T-schema is preserved. This is quite an achievement, as the T-schema threatens the consistency of accounts which allow gappiness. Assume that neither A nor $\neg A$ is true. Then by the T-schema and *modus tollens*, both $\neg A$ and $\neg\neg A$ are true, which is a contradiction. According to supervaluationists, this little argument contains a fallacy of equivocation. However, it isn't clear how non-supervaluational accounts are to avoid it.

In general I suppose the move will be to render the T-schema impotent, as van Fraassen (1966) does in his attempt to defend weak supervaluationism. He claims that we lose the T-schema, but this is no great loss as both " $A \vdash 'A' \text{ is true}$ " and " $'A' \text{ is true} \vdash A$ " are still valid. However, his argument for saying these are valid relies on there being no higher-order vagueness, and on validity being mere truth-preservation, rather than the stronger rule outlined at the end of the last section. Each of these premises seems to be mistaken.

The argument for supervaluationism here relies, it might be thought, a little too heavily on the weaknesses of its rivals. In part that's not accurate, some of the objections above are to all possible opposing accounts, not just those now on the market. However, in part it's an unavoidable problem. The challenge Fodor and Lepore put forward can be easily met for one class of sentences (the first-order sentences), the difficulty is just expanding it beyond that. Since the main ground for expansion is theoretical simplicity and the coherence of a unified approach, much of the discussion is going to be concerned with counting the costs of going this way as compared to going another. Once that point is reached, negative arguments seem likely to be most prominent.

Fodor and Lepore do anticipate somewhat this general line of reply. In a long footnote, they note that some critics have suggested the conceptual truths they want precisifications to honour are

metalinguistic, higher-level truths. They first suggest this oughtn't matter. I take it this is the same as pressing the question of why, given our justification, we should apply supervaluational semantics to higher-order sentences. This objection I've already considered. They then claim the truths are not higher-level with the following argument.

The crucial consideration is that you cannot make a man more (/less) bald without altering his hair-to-head ratio. So if there is any valuation on which ['Al is bald'] is true (/false /indeterminate) and Al's head-to-hair ratio is m/n , then ['Al is bald'] is true (/false /indeterminate) in every valuation in which Al's head-to-hair ratio is m/n . (523n)

The question turns on whether (1) is true.

(1) If A is true on any precisification, it is true on all of them.

Surprisingly, there is an argument for (1), just the little argument from facts about supervenience that Fodor and Lepore give. I take it this can be generalised from sentences about baldness to sentences generally. The moderate supervaluationist respects this argument; on their view (1) is ambiguous. The standard reading would have it come out at least possibly false. However, there is a reading which we get by prefixing 'According to all precisifications' to (1), on which it is necessarily true. The simple fact which this reading expresses is that according to a precisification, it is the only acceptable precisification. However, trying to use that to show there can't be more than one precisification, or that metalinguistic sentences aren't of a different type to ordinary sentences, rests on a fallacy of ambiguity.

5.7 Local and General Supervaluationism

I think a broad supervaluational program succeeds; supervaluational semantics are correct for at least sentences with vague predicates, and possibly also for ambiguous sentences. However, even

if this program were to fail, I think a supervaluational program for probability sentences would still succeed, because the objections to supervaluationism seem particularly weak when applied here.

It is conceded that the precisifications that supervaluationists use when interpreting vague terms are fictions. All of these precisifications have some features which differ from English; they are all precise. On the other hand the precisifications we use for interpreting probability sentences are much more natural. Each of them does represent a reasonable epistemic state, even though, since all of them are precise, there are reasonable epistemic states which are not precisifications. So the differences which Fodor and Lepore stress between models of languages and those languages are not real differences here.

There is also a distinction in terms of theoretical priority. In a natural language the vague terms are theoretically prior; we introduce the precisifications as a theoretical device for understanding these terms. In this case we start at a more theoretical position. The only contribution of natural language is to draw the link between probability and reasonable credences. What counts as reasonable will itself be largely determined by theoretical reflection. As it happens, the theory I have used here (developed in chapter 3) takes the precise probability functions as the basic elements, and allows degrees of belief to be imprecise by allowing them to be vague over a set of these functions. In these cases, it isn't too surprising that we should define truth in terms of what is true according to each of these reasonable functions.

To put this point differently: in understanding probability, as in understanding credences, I take comparative, qualitative sentences to be primary. A probability function is reasonable iff it makes all of the comparative sentences true. Given that the comparative sentences are primary, and the admissible probability functions make all of them true, it is no great stretch to say that true probability sentences are just those true according to all functions. If we didn't have the theoretical

arguments for the importance of probability functions (as opposed to just any old functions which made the comparative sentences true) this move would look quite unjustified; but with those arguments, it looks a little bland.

So in sum, even if the defences constructed above crack and supervaluational approaches are shown not to work for interpreting vague sentences, the specific usage of supervaluational semantics here may still be secure. I hope that doesn't happen. I would like to see a general supervaluational program succeed, but I don't think the fate of my theory of probability depends upon it.

5.8 The Reasonableness of Imprecision

It is implicitly assumed above that there is more than one reasonable probability function. Some arguments for this were mentioned in chapter 3. This section will examine those arguments in more detail. There are three important classes of arguments for that conclusion: arguments from ignorance; arguments from rational disagreement and arguments from vagueness.

5.8.1 *Ignorance*

There are some propositions about which we have next to no evidence. Classically the probability of these propositions was determined by Laplace's principle of indifference. So, if we know, for example, nothing about a coin that's about to be tossed, we ought assign equal credence to the proposition that it will fall heads if tossed as to the proposition that it will fall tails. In general, the principle says that we can infer from the absence of evidence about p and q that the probabilities of the two are equal.

In gaming situations this does not lead to implausible results, but when applied more widely it can do so. The classic examples of this are first presented in Bertrand (1889). This example is first due to von Kries (1886). A factory makes cubes of random sizes, the largest having a side

length of 2cm, and the smallest having a side length of 0. What is the probability that the last cube to come off the production line had a side length greater than 1cm on just the evidence we have? By applying Laplace's principle we seem to get the answer $1/2$; the evidence that the cube's side length is larger than 1cm is equal to the evidence that its side length is less than 1cm. However, we could instead look at the evidence about what the volume of the cube is. We know that the volume is at most 8cm^3 , and for every interval $[n, n+1]$ ($0 \leq n \leq 7$) the evidence that the volume is in that interval is identical. So the probability the volume is less than 1cm^3 is $1/8$. But the volume is less than 1cm^3 iff the side length is less than 1cm. Hence the probabilities of the two should be the same. It seems that Laplace's principle leads to inconsistency.

The importance of this is that, without Laplace's principle, there is no way to represent ignorance about p through a single probability function¹⁵. Since ignorance is, it seems, possible, we should be able to represent it. And we should be able to say what epistemic states are reasonable under ignorance. If we allow multiple reasonable probability functions, this is possible, as has been shown. This argument has some historical importance, as it seems it is what convinced Keynes (1921a: ch. 4) to allow degrees of belief to be incommensurable.

There are several related arguments to this one which are worth stating, but which have little independent force. In traditional epistemology, an agent is allowed to not have beliefs about a proposition. That is, we distinguish ' A believes $\neg p$ ' from ' A doesn't believe p '. In natural language words which would seem to indicate the latter often are used for the former, but this doesn't mean the distinction isn't real.

Many probabilistic epistemologies are not so liberal. So traditional subjectivism demands that an agent have a precise credence in a proposition, and that this be 1 less their credence in that

¹⁵ This conclusion might be a bit quick; the constructivist theory developed in later chapters promises to do just this without Laplace's principle. But it is clearly right for orthodox probability functions.

proposition's negation. In other words, it demands the agent have an epistemic attitude towards that proposition, a demand from which traditional epistemology recoils. The epistemology here is probabilistic but it allows reasonable agents to take no epistemic stance towards propositions. If an agent has no thoughts at all about a proposition p , then their epistemic state will be represented by a set of probability functions which includes a function Pr such that $Pr(p) = x$ for all x in $[0, 1]$. An argument similar to this one is used in Jeffrey (1983) for allowing probabilities to be vague.

5.8.2 *Rational Disagreement*

I take it as a datum that reasonable people can, on the same evidence, have different degrees of belief in the same proposition. That is, there is more than one reasonable response to certain bodies of evidence. I don't, for example, regard it as evidence that one or other discussant is unreasonable, if participants in a debate about the Kennedy assassination express different credences in Oswald's guilt. The evidence permits some doubt; how strongly that doubt is felt seems to vary among reasonable people.¹⁶

In other words, I take it that the possibility for reasonable people to differ in the likelihood of a hypothesis given some evidence is something that a theory of probability should be required to explain. If there were only one reasonable probability function this explanation would be hard to find. However, on the theory adopted here it is easily explained because of the existence of multiple reasonable probability functions.

¹⁶ This is a little speculative, for 'reasonable' here entails 'coherent', and few people's beliefs are fully coherent. So we can't verify this claim by simply finding reasonable people with the same evidence who disagree. However, it seems implausible to think the kinds of disagreement mentioned in the text would dissipate were the disputants to obtain perfect computational powers.

Some people may dispute the intuition on which this argument relies. If so I have no response; I hope the other arguments are more convincing. A more important objection to this argument is that it looks like an *ad hoc* move. The possibility of rational disagreement seems to threaten the whole project of analysing probability as reasonable degree of belief. It is incorporated only by weakening the semantics to allow for imprecise probability statement. If this argument were the only ground for allowing multiple reasonable probability functions, the objection would I think succeed. But that isn't the argument I am using here. Like Keynes, I take the need to represent epistemic states with minimal evidence as the primary ground for allowing imprecision. That the most natural way to do so has this pleasant side-effect, that it allows rational disagreement, should count as evidence we're on the right track.

5.8.3 *Vagueness*

There is a more direct argument for allowing imprecision. Try to determine, to the fifteenth decimal place, your credence that the Bulls will win the next N.B.A. title, or the Democrats the next Presidential election. It simply can't be done; credences are not that precise. We have to have some way of representing vagueness, and the method used here seems better than the alternatives.

This argument is a bit quick. It might be responded that we do really have precise credences but they are subject to ordinary measurement errors. This is the response Borel (1924) makes to an imaginary objection like the one above. It might be, as Williamson (1994) suggests, that vagueness is in general an epistemic phenomenon; there is a sharp divide between the tall and the not-tall but we are unaware of what it is. Similarly there might be precise credences we have in all propositions of which we are unaware.

Borel's response would work if there were independent reasons for thinking credences are precise. We could then explain away the anomalous introspective evidence as a measurement problem. But I don't see what those independent reasons could be, and without them the most natural

interpretation of the evidence, that credences appear imprecise because they really are, seems most appropriate.

Whatever the general merits of Williamson's resolution of the problem of vagueness that type of approach seems inappropriate here¹⁷. There is nothing implausible about insisting we can have (partial) beliefs of which we are unaware. So we might have a precise credence in the Bulls winning but be unaware of it. On a broadly functionalist view, provided we have all the right dispositions that go along with believing the Bulls will win to degree 0.4, we do believe that to degree 0.4, whatever the introspective evidence. It is, however, just as implausible that we have sufficient dispositions to settle our credence to the fifteenth decimal place as that we could determine that credence introspectively.

It might be thought that by considering our dispositions to accept or reject bets at certain odds we could settle precisely our credence in a given proposition. Apart from the general objections in earlier chapters there is a particular problem here. Assuming we do have dispositions to accept or reject any bet is quite implausible. For some bets we simply don't have dispositions about what to do when faced with them. We might, on reflection, choose to accept (or reject) them, but this reflection may involve changing our epistemic state. This means that even if an agent does accept a bet after being offered it, she needn't have originally had a disposition to accept that bet. Any defence of the claim that all degrees of belief (or even all reasonable degrees of belief) are precise will turn out to rely on an implausible claim somewhere, so this direct argument for allowing imprecision seems to be quite strong.

¹⁷ Williamson hasn't written about the imprecision in degrees of belief so I don't know if his general solution is intended to apply here.

5.9 The Reasonableness of Precision

I have made two assumptions above which might be considered controversial. The first is that it is always reasonable for an agent to believe any proposition to a precise degree. The second is that the union of all reasonable probability functions represents a reasonable epistemic state. The purpose of this section is to defend these assumptions. For simplicity in this section I'll use the phrase 'epistemic state' to refer to a set of probability functions representing an epistemic state.

I'll start with the defence of the second assumption. I don't want to defend the claim that whenever C_1 and C_2 are reasonable states then $C_1 \cup C_2$ is a reasonable state. I think reasonable states must satisfy some kind of continuity principle, so that the agent's degree of belief in a proposition must be an interval. I do, however, want to defend the related claim that whenever C_1 and C_2 are reasonable states then some superset of $C_1 \cup C_2$ is a reasonable state. This just amounts to the claim that whenever two reasonable agents disagree on the probability of some proposition, there is a reasonable state which is neutral on the question of which of them is right. That seems plausible enough; if all participants in a debate are being perfectly reasonable, there is no requirement on a reasonable agent that they make a decision between the participants. From this principle it follows straightforwardly that some superset of the union of reasonable epistemic states is a reasonable epistemic state.

The first assumption is a little harder to defend. Some writers have thought that the arguments of the last section show more than I've intended. They show that under some circumstances, particularly when there is little evidence about a proposition, it is unreasonable to believe that proposition to a precise degree. This, for example, seems to be the view Keynes takes. On any evidence set there is only one reasonable epistemic state, and subsets of that state are not reasonable. The position adopted here is that all subsets of reasonable epistemic states are reasonable.

The first argument for this relies on conclusions not yet justified. In subsequent chapters I will be arguing for a decision theory called Caprice. According to this theory, a set of choices and dispositions to choose is reasonable iff the agent is reasonable and there is a probability function P_i in the agent's epistemic state such that, for every choice of (or disposition to choose) A over B , the expected value of A according to P_i is at least as great as that of B according to P_i . There are no restrictions on how an agent chooses P_i from her epistemic state¹⁸. So an agent whose epistemic state is C , where $P_i \in C$, will make exactly the same decisions as an agent whose epistemic state is $\{P_i\}$. If this decision theory is right, it follows that precisifying (i.e. an agent moving from epistemic state C to epistemic state C_1 , a subset of C) can never lead to the agent making decisions which would have been irrational according to the coarser epistemic state. So from a purely pragmatic perspective, there is no cost to precisifying, hence it is not irrational.

There is more to being reasonable than not losing money, so I'm sceptical about the force of this argument. A stronger argument turns on the possibility of rational disagreement. Assume epistemic state C is the largest reasonable state. There are three possibilities: all subsets of C are reasonable; some but not all subsets of C are reasonable, and finally that no subsets of C are reasonable. We are wanting to show the first is correct, so assume for now it is false. The second seems implausibly arbitrary; what ground could there be for distinguishing the reasonable precisifying moves from the unreasonable ones? And the third does not allow reasonable agents to disagree, it requires all reasonable agents to have the same (imprecise) degree of belief in a proposition on a given evidence set. This is the position Keynes adopted, but it is refuted by the possibility of reasonable disagreement. So the first option, which is what I wanted to defend, is the only one left standing.

¹⁸ This language of choosing P_i is perhaps a bit misleading. The picture I have in mind is that the agent simply makes choices and if she is reasonable the condition will sort itself out. There is no conscious decision to choose according to a particular probability function.

Chapter 6

Objections

6.1 Introduction

So far I have defended two distinctive views about probability. The first is that probability sentences, properly construed, are non-contingent. The second is that the probability relation between propositions is not always numerical. In chapter 5, I relied on the technique of supervaluations to explain how these non-numerical relations interact. I call any analysis of probability which defends the first of these theses a *necessitarian* analysis, and any analysis which defends the second an *imprecise* analysis. (The imprecision in question is in the values of the probability relations, not in the analysis!) The first writer to defend each of these theses was Keynes (1921a). Several of the objections that I'll discuss in this chapter were first made by Ramsey and aimed at Keynes's theory, so before commencing I need to note one important difference between my theory and his. Keynes, at least in his *Treatise on Probability* held a *logical* analysis of probability. Some of the objections to Keynes I'll discuss are objections to the logical analysis, but are not objections to a necessitarian view. To defend my theses, I don't need to respond to these. Indeed, as I pointed out in section 1.7, it is because I accept some of them that I reject the logical analysis.

One other distinct feature of Keynes's theory should be noted. I mentioned in chapter 5 that we could think of the probability of p given q as a function from distributions to elements of $[0, 1]$. Keynes says that probabilities can be non-numerical, but he doesn't give this analysis of their 'internal structure'. He is happy to say the probability of p given q is α , where we might know no more about α than $\alpha > 0.2$ and $\alpha < 0.6$. By means of some axioms, he shows how we can interpret addition and multiplication of these probability-values, but it is rather unclear whether or not addition and multiplication still have the same meaning they do in natural language once these axioms are added. As a consequence, when I claim that objections of Ramsey's that are directed

at Keynes's refusal to discuss the structure of probability relations don't harm my theory, I am not attacking Ramsey (his objections usually hit their intended target) but merely showing why my theory isn't doomed because of his work.

Since Ramsey levelled so many distinct objections to the Keynesian project in his 1926 paper *Truth and Probability* I will spend most of this chapter analysing and where necessary responding to those objections. It is commonly assumed, particularly in non-philosophical discussions of Keynes's work¹, that the Keynesian theory was fatally wounded by Ramsey's early attacks on it. I will argue that while Ramsey's attacks seem to work against logical theories, they aren't overly persuasive against necessitarian views. I will argue in chapter 10 that Keynes held a similar, if less worked out, view of the value of Ramsey's attacks. There are several separate arguments against Keynes squeezed into Ramsey's paper. For ease of later reference, I'll first simply list what I take to be the main arguments before attempting any kind of response.

- There aren't any probability relations of the kind Keynes requires.
- If there are such relations, we can't determine what their value is when the relata are simple.
- It is mysterious how the probability of p can go from being incomparable with any given number to being numerically precise by the addition of evidence.
- It is unclear why Keynes's probability relations should obey the laws of the probability calculus.
- Keynes's theory relies on the discredited Principle of Indifference.
- Keynes's theory requires that our evidence be known for certain, but much of the time our evidence is vague and uncertain.

Unlike Keynes, Carnap backed up his necessitarian theory with a detailed calculus for the probability logic. For various reasons, this calculus has attracted more attention than the

¹ See, for example, Bateman (1996) and Runde (1994a). From a more philosophical perspective, the same assertion is made in Zabell (1991: 224).

philosophy underlying it. Howson and Urbach (1989) provides one exception, and I deal with their objection to Carnap at the end. I suspect their worry is widely shared, and their presentation sufficiently representative, so my replies may be relevant to many critics.²

6.2 There are no such things

But let us now return to a more fundamental criticism of Mr. Keynes' views, which is the obvious one that there really do not seem to be any such things as the probability relations he describes. He supposes that, at any rate in certain cases, they can be perceived; but speaking for myself I feel confident that this is not true. (161)

In his original theory, Keynes held that we could perceive directly the probability relations between propositions, in the same way that we perceive entailment relations. Some have argued that this paragraph was a decisive refutation of Keynes's Platonist assumptions. It is, of course, nothing of the sort. Ramsey does not deny that we can perceive, in some sense of that word, entailment relations. Nor does he deny that we can perceive relations, such as having a common subject or predicate, which hold between propositions, at least on a Tractarian view of what propositions are. All he rejects is that we can perceive the particular relations Keynes posits.

From the way we have put forward the theory, this is a rather odd thing to say. He couldn't be saying that he doesn't perceive that there are some pairs of coherent belief states and evidence sets which are reasonable and some which are unreasonable. In fact he goes on to say that a person who doesn't draw inductive conclusions from their evidence is unreasonable (197). However, he doesn't believe that such an 'unreasonable' person would 'sin against formal logic or formal probability'. Now this may be so on a narrow construal of 'formal'. If, however, we are internalist about epistemic justification then we are committed to saying that we can perceive that such a

² All page references in this chapter unless otherwise stated are to Ramsey (1926a).

person's epistemic states are *necessarily* unreasonable³. So, to perceive that such a person is unreasonable is to make a perception which, if it is correct, is necessarily correct.

On the theory advocated here, if for every value of n we can perceive whether believing p to degree n on a certain body of evidence q is reasonable or unreasonable, then we can perceive the probability relation between p and q . We won't ordinarily be able to do this because the boundaries between the reasonable and the unreasonable will be vague, but this merely corresponds to the probability relation being vague. The crucial point is that there is nothing more (or less) to a probability relation than a bundle of facts of the form *It is (not) reasonable on such-and-such evidence to have this degree of belief in p* . Since Ramsey is happy to say we can perceive facts of this latter sort, it follows that we can perceive probability relations.

Ramsey goes on to say that, because other people can't agree on the value of probability relations, he believes no one else perceives them either. In part this criticism is met by the theory of vague probabilities set out in chapter 5. There it was argued that reasonable people could have different degrees of belief in the same proposition on the same evidence. However, in these cases, the probability of a proposition on some evidence isn't *the* degree of belief a reasonable person would have in the proposition on that evidence. There is no such degree, even if we allow non-numerical degrees of belief. Rather there are a range of reasonable degrees, and the value of the probability relation is this set.

6.3 Probability Relations Between Simple Propositions

If, on the other hand, we take the simplest possible pairs of propositions such as 'This is red' and 'That is blue' or 'This is red' and 'That is red', whose logical relations should surely be easiest to see, no one, I think,

³ This will be false unless we take 'epistemic states' to be a rigid designator, which for ease of exposition I do.

pretends to be sure what is the probability relation which connects them.

(162)

Ramsey might concede that we can work out the probability relation between two complex propositions such as our entire current evidence and “Oswald killed JFK”. However, as the above quote indicates, we can’t tell what the probability relation is between simple propositions. This is entirely at odds with the rest of logic, where we are prepared to rely on agreement about the relationships between simple propositions to work out the relationships between complex propositions.

The easiest way to shrug off Ramsey’s objection here would be to say that it only directly attacks logical analyses of probability rather than necessitarian analyses. However, the point deserves some more discussion. After all, even if these examples aren’t central on a necessitarian view, as they are on a logical view, we are still committed to saying that there is a probability relation between them.

In previous chapters I have had occasion to identify probability relations with the set of values they take. This is a convenient shorthand, though it does lose some information. In this notation, I can say quite precisely what the conditional probability of *This is red* given *That is blue* is. It is the interval $[0, 1]$. In other words, there is no degree of belief in *This is red* which is either ruled out as irrational or ruled in as the only rational response on this minimal evidence. This response to Ramsey doesn’t entirely succeed, because the interval notation isn’t fully informative, but it is a start.

The other response we can make to this objection is that Ramsey’s assessment of simplicity is very much related to his Tractarian conception of propositions. This analysis of propositions is crucial to the logic of decision he develops, and for which the paper is more well-known. If we analyse propositions as sets of possible worlds, and assume possible worlds are the ‘ultimate

organic unities' (177) which are the subject of choices, we cannot use Ramsey's method for determining beliefs and desires. On my preferred analysis of propositions, *This is blue* will not be a simple proposition in any interesting sense. On the contrary, it is as heterogenous a set of possible worlds as one could care to imagine. The truly simple propositions will be those which are true at one world only. However, the linguistic representation of such propositions will be infinitely complex. Hence *simplicity* is theory-dependent, and even if I agreed with Ramsey that not being able to say what the probability relation was between simple propositions was problematic, these examples wouldn't count against my theory.

6.4 The Sorites Objection

[I]t is hard to suppose that as we accumulate instances there is suddenly a point, say after 233 instances, at which the probability relation becomes finite and so comparable with some numerical relations. (162)

On Keynes's view, as on mine, when we have observed a billion red round things, and no non-red round things, the probability that *a* is red given that *a* is round is greater than 0.99. That is, it enters into numerical comparison. This is necessary if we are to have any projectible predicates. However when we reduce the number of observations from a billion to, say, one, the resultant probability is not comparable with any numbers, or numerical relations as Ramsey puts it. Ramsey seems to think this combination is implausible, because it would require an arbitrary number of observations after which the probability does become numerically comparable.

It is a remarkably weak objection, particularly by Ramsey's standards. Assume a certain pile of sand is not a heap. I can, by adding one grain at a time, end up with a heap of sand. However, I don't have to assume there is some point, say 233 grains, at which it suddenly becomes a heap. Given this it is hard to see what the basis of Ramsey's objection is. As an attempt to give some bite to the objection in the previous section it doesn't seem particularly plausible.

6.5 The Probability Calculus

For now it is easily seen that if partial beliefs are consistent they will obey these axioms, but it is utterly obscure why Mr. Keynes' mysterious logical relations should obey them. (188-9)

As I noted in the introduction to this chapter, for Keynes the probability relation between two propositions may be some non-numerical value. There are then two related objections that can be made concerning the calculus of these values. The first is that it isn't at all clear what we mean when we add or multiply them. For example, if the probability of p given q is α , and the probability of $\neg p$ given q is β , Keynes says we can conclude $\alpha + \beta = 1$, but he doesn't say what this might mean. The second objection is that it isn't clear why he should want it to be the case that $\alpha + \beta = 1$. I take the results of chapter 3 to provide an answer to this question, though to be fair to Ramsey it is an answer which appears nowhere in Keynes.

So the more important challenge is to the meaningfulness of the mathematical notation in Keynes's theory. If \mathbf{P} is the set of reasonable probability functions, we can identify the probability of p given q with a function from \mathbf{P} to $[0, 1]$. In this notation we'll have to identify numbers with constant functions, i.e. identify 1 with the function, call it $\mathbf{1}$, such that $\mathbf{1}(Pr) = 1$ for all Pr . Now we have clear concepts of what it is to add and multiply functions, at least with common domains. For functions f, g and h , we can easily say $f + g = h$ iff for all $x, f(x) + g(x) = h(x)$. It seems plausible enough to say the '+' in the first equation means the same as the '+' in the second⁴. Similar definitions can be given for the multiplication and division of functions.

Now, with this notation, we can see how it is possible that $\alpha + \beta = \mathbf{1}$ (which we are identifying with 1) how some probability relations can be non-numerical (because they are not constant

⁴ Though perhaps ever since Wittgenstein (1953) we ought be a little sceptical about such claims of meaning extension. I hope everyone will agree that I have used '+' in an acceptable way, and if Wittgenstein is right that will mean I have used it acceptably.

functions), and how this doesn't prevent us applying addition and multiplication operations. Thus I presume the challenge as to the meaningfulness of Keynes's algebra is met, and its justification is given by chapter 3.

6.6 The Principle of Indifference

Secondly, the Principle of Indifference can now be altogether dispensed with; we do not regard it as belonging to formal logic to say what should be a man's expectations of drawing a white or a black ball from an urn; his original expectations may within the limits of consistency be any he likes; all we have to point out is that if he has certain expectations he is bound in consistency to have certain others. This is simply bringing probability into line with ordinary formal logic, which does not criticise premises but merely declares that certain conclusions are the only ones consistent with them. (189)

This objection isn't applicable to my theory, but it is worth considering as an objection to logical analyses of probability. A logical analysis of probability has to elevate the Principle of Indifference to a logical theorem. This might be thought problematic because of the various paradoxes of indifference discussed in 5.8.1. Even if I can avoid the paradoxes somehow, and as I noted there allowing probabilities to be vague seems to do the trick, there is a lasting impression that the ideal solution would have been to not allow the problems to arise. So I suspect the intuition Ramsey has, that if we can do away with the Principle we ought, is just the right one to have.

It may be possible to develop a precise logical analysis of probability which is immune to all of the paradoxes, but given the calibre of the theorists who have tried and failed, it seems doubtful. It would be more plausible to think an imprecise logical analysis which avoided the paradoxes could be developed. The two grounds I have for thinking this won't be done are, again, that we

would expect if it can be done it would have been done already, and that there seems to be a fairly mechanical procedure for constructing objections to any approach.

As it turns out, the formal logic of probability defended here is just the same as that Ramsey defended, with the only possible exception being that we might regard the Principal Principle as a logical rule. Nevertheless, when we say someone has made a logical error in their allocation of degrees of belief, we are just saying that their beliefs are inconsistent. In Ramsey's terms, the person making an error has failed to do something they are 'bound in consistency' to do. As Carnap points out (1950: 337), the formal component of all theories prior to his shared a common logic. So even though Keynes had argued that there were logical restrictions on what degrees of belief people could reasonably have in propositions, when he wrote the formal component of his probability logic these restrictions were not incorporated.

There is one infelicity of expression in the Ramsey quote above. When Ramsey says that a man's expectations 'may within the limits of consistency be any he likes', the modal *may* is being used rather oddly. He doesn't mean that having any old expectation would be epistemically acceptable all things considered. Rather he means that any expectation which is consistent is *logically* acceptable. As he goes on to say in his discussion of induction, there are restrictions on what is epistemically acceptable, i.e. on what is reasonable, which are not logical restrictions.

Finally, we should note the oddity in the last sentence of the quote from Ramsey. It might not be part of logic to criticise the premises that various people hold, but it is part of epistemology broadly speaking. All we ought to conclude from what Ramsey says is that the logic of probability should have little to say with regard to criticising individual probabilistic judgements. I agree; the theory defended here does not make *logical* criticisms of agent's whose beliefs are coherent. However, the boundaries of critical epistemology are not the boundaries of logic. This is true in non-probabilistic epistemology, and it is true in probabilistic epistemology. In sum, I agree with

Ramsey that logical analyses of probability rely too heavily on a Principle of Indifference; but I disagree with his claim that we have no grounds for ruling as unreasonable any consistent beliefs. It is this last possibility which opens up an necessitarian theory of probability.

6.7 Uncertain Evidence

I think I perceive or remember something but am not sure; this would seem to give me some ground for believing it, contrary to Mr Keynes' theory, by which the degree of belief in it which it would be rational for me to have is that given by the probability relation between the proposition in question and the things I know for certain. (190)

There are two responses we can make to this objection. The first is that it doesn't apply on my conception of evidence; the second is that this is a problem for the application of the theory, rather than the theory itself. Ramsey's observation is of undoubted importance for anyone wishing to construct a machine which has reasonable beliefs, or if you doubt machines can have beliefs, functional states which behave like beliefs and would be reasonable were they believed. However, it is unclear why such practical worries should harm the theory under construction here.

On the view endorsed here, the evidence in a probability relation is a set of worlds in all of which I have certain experiences. The worlds can differ from the actual world in any way at all, as long as my experiences are held constant. The upshot of this is that my evidence isn't of the form *There is a blue book on my desk*, but rather *I am observing a blue book on my desk*. Even the latter isn't a precise representation of the evidence, since my experiences might be being caused by an evil-demon, but for practical purposes it is as close as we can get in language. In cases considered by Ramsey, the evidence wouldn't be of the form *It's probable I saw a blue book on my desk last night*. Rather, it would be of the form *I am having a dim memory of seeing a blue book on my desk*. The reasonable degree of belief in the proposition that there was a blue book on my desk

might be moderately high on the basis of the evidence, but the evidence is still something that is taken to be certain.

The other response to make here is that the concept of ‘reasonableness’ I am using requires, at least in its technical aspects, superhuman ability. After all, it is unreasonable on this picture not to realise what can be entailed from what, not an ability many of us possess. Hence the question of how we should incorporate failing memories into the theory is of the same type as the question of how we should incorporate failing inferential processors. That is, not a relevant question at this level of abstraction. Rather it is something that need to be incorporated in less abstract theories. Is it legitimate to brush aside these concerns as, in effect, engineering problems? It is, because when we are designing practical systems to approximate ideally rational systems, we have to know what it is we’re approximating. The purpose of our theoretical pursuits, *vis a vis* such projects, is to work out where the goal posts are. Once there is agreement on what the theoretical aims of a practical rational system are, we can assess how well the system achieves those aims. However, without this agreement, we can’t evaluate such systems. Given then that our aim is to investigate the ideal, it might be plausible to abstract away from the difficulties such as failure of memory that worry Ramsey.

Kaplan (1996: 36-8) suggests a useful way to understand the demands that theories like this one are making. When we say that, for example, epistemic states ought be consistent, we are not saying that it is a legitimate criticism of a believer that their belief states are inconsistent. However, it is a legitimate criticism of their belief states, something the believer will usually agree with, as they’ll attempt to remove inconsistencies brought to their attention. Similarly in the case of memory, it is no legitimate criticism of an agent that their memory is less than perfect, however it might well be a legitimate criticism of some states of that agent. And again the agent will often agree in the sense that they’ll say it is best, other things being equal, to remember more of our evidence rather than less. However, here other things are never equal, there are always trade offs

to be made, and the agent is only subject to legitimate criticism if in making this trade off they don't take the ideal of perfect evidence retention seriously enough.

6.8 A Recent Addition - Dependence on A *Priori* Assumptions

Most critics of Carnap's theory of probability have focussed on the technical aspects of his theory. There are a few good reasons for doing this. First, his technical theory is interesting for its own sake. Secondly, if there are major flaws in the particular technical theory (as seems to be the case) this counts against the general approach, both because there is arguably a burden on Carnap and his followers to produce a sound technical theory, and more generally because we might suspect that if Carnap couldn't complete this project it can't be completed. However, it would be nice to see in the literature more discussion of the philosophy behind Carnap's theory, and solid objections to it. As my necessitarian approach adopts Carnap's philosophy, but not his technicalities, I incur a duty to answer those objections directed at that philosophy.

The most serious objection is that Carnap's method is too *a prioristic*. The discussion in Howson and Urbach (1989: 52-56) seems to be the most substantial presentation of this objection. I say it is necessary, and arguably *a priori*, which functions Pr are reasonable. For example, let my evidence proposition be E , and let q be the proposition that the moon is made of green cheese. Then I say it is necessary that $Pr(q | E) < 0.9$ for all reasonable Pr . Note that every Pr is, as well as being a conditional probability distribution, an unconditional probability distribution, since we can define $Pr(A)$ as $Pr(A | T)$. Hence I'm committed to there being *a priori* probabilities. Howson and Urbach object.

For any conditional probability distribution over the sentences of a language necessarily involves the assignment of unconditional probabilities to a partition of the space of possibilities representable

within L . But what considerations can possibly justify any such *a priori* distribution? (Howson and Urbach 1989: 53-4)

Any *a priori* probability distribution ... is going to be arbitrary. For this reason we do not regard people who try to evaluate the probabilities of hypotheses relative to data as doing exercises in a genuine logic of generalised deduction, for we take logic to be essentially noncommittal on matters of fact. (Howson and Urbach 1989: 55)

Howson and Urbach back up these claims by pointing out, rightly, that considerations of symmetry or simplicity will not give us the *a priori* distributions. Any time we try to make the distributions symmetrical or simple relative to one set of considerations we will make them more skewed and more complex relative to others. I accept this, but not the conclusion they draw from it.

Three relatively inessential points before we start. First, in my story (unlike Carnap's), probabilities are assigned to pairs of propositions, not pairs of sentences, and propositions are just sets of possible worlds. These 'possibilities' are in general not representable within any language. This I take it is no response at all to their objection. Secondly, I *could* avoid making my theory *a prioristic* if I let the reasonable probability functions be whatever play a certain role in the actual world. That is, I could consistently with my theory say that which probability distributions are reasonable is necessary *a posteriori*. As I noted in section 4.4, the relevant simplicity and symmetry considerations might be given by our actual practices. However, that approach has problems, and to adopt it just to avoid a charge of *a priorism* would be untenable. In any case, Howson and Urbach could rewrite their objection to deal with this. So I'll write here as if necessary and *a priori* were interchangeable.

Thirdly, and this is a bit important, it isn't true on my theory that a conditional probability distribution 'necessarily involves the assignment of unconditional probabilities' *a priori*. Let an epistemic state be represented by the set $\{Pr: Pr(q | E) = 0.2\}$. One conditional probability is quite precisely defined here, but no unconditional probabilities of contingent propositions are defined to be sharper than $[0, 1]$. I suspect the reference to unconditional probabilities was more a rhetorical flourish than a crucial part of the argument, which is why I think this point is mostly unimportant.

Those clarifications aside, I can proceed. I trust the reader agrees that having degree of belief 0.9 or higher in q on evidence E is unreasonable. If not, please reconsider. If so, read on. Following Lewis (1980) I have defined reasonable probability functions to be those which licence no unreasonable degrees of belief. So we can conclude that all reasonable probability functions Pr are such that $Pr(q | E) < 0.9$, because this just means that having degree of belief 0.9 or higher in q on evidence E is unreasonable. Call this conclusion F . Question: Is F empirical or *a priori*?

If F is *a priori*, then I have a response to Howson and Urbach, for I can say that what determines the set of reasonable probability functions is just the set of *a priori* facts like F . Note that their claim that logic is noncommittal on matters of fact is clearly mistaken unless we say *a priori* truths are not 'matters of fact'. This is perhaps a non-standard use of fact, but not an unintelligible one. However, on it, F does not turn out to be a matter of fact, so my 'logic' is noncommittal on matters of fact.

Hence for their objection to work, F must be empirical. But what evidence could there possibly be for F ? Or perhaps more strikingly, what evidence could there be for $\neg F$? What is one meant to say about my counterpart who has had exactly the same experiences as I, but believes to degree 0.9 that the moon is made of green cheese, and does so reasonably. I very much doubt I have such a counterpart. This does not show that F is not empirical, it might just show that E entails F . If

that were true, it might be the case that F is empirical but nevertheless I could have no such counterpart. The problem with that is that it is inconsistent with the conclusions about updating I derived in chapter 3. That is, it is inconsistent with my insistence on Conditionalisation. In fact, it is inconsistent with virtually any interpretation of conditional probability, as this example shows.

Let E_1 be strictly weaker than E . That is, E entails E_1 but not *vice versa*. And let E_1 be such that it doesn't entail F . Now say I have a counterpart whose evidence is E_1 . Perhaps it would be reasonable for him to believe q to degree 0.9. Since he doesn't know F , his evidence is insufficient to support it, and for all I've said F might be false in his world. So, for all he knows, it is reasonable to have probability functions such that $Pr(q | E) = 0.9$ in the representation of his epistemic state. In fact, if I'm right and any precisification of a reasonable epistemic state is reasonable, he could reasonably have his degrees of belief represented by just that function. But then were he to learn E , there would be nothing he could reasonably do. By conditionalisation, he would have to have degree of belief 0.9 in q . However, he would now know F , so he would know that his representative probability function was unreasonable. Hence, he would know that the only thing he can reasonably do – i.e. conditionalise – is unreasonable. So his epistemic state provides no consistent guidance in a possibility he envisages as possible. But, by definition, no reasonable state does this. Hence, his state was unreasonable to start with; so, having a function $Pr(q | E)$ in one's representor is always unreasonable. We can rephrase all this without the assumption that precisifications of reasonable degrees are reasonable, but it's less clear. The point again is that, when my counterpart conditionalises on evidence E , he won't be able to completely precisify.

So my response to Howson and Urbach is in two parts. First, it is obvious we do believe in F and like claims. Secondly, it is implausible to say that F is empirical. So, that necessitarian theories are committed to *a priori* assumptions like F is no mark against them, and may in fact be a benefit.

Chapter 7

Philosophical Predecessors

As has already been mentioned, several authors have argued that imprecise or non-numerical degrees of belief ought to be permissible. The aim of this chapter is to look at four recent exponents of this view, and in particular at their motivations for allowing imprecision and the technical frameworks they develop to deal with them. My discussion of the writer to whom my position is closest, Keynes, will wait to a separate chapter investigating the connections between Keynes's theory of probability and his economic theories.

7.1 Levi

In several papers and books, Isaac Levi has developed a theory for reasoning and decision making which allows that degrees of belief, what he calls *credal probabilities*, can be imprecise. He also allows that values can conflict, which might be captured by allowing utilities to be imprecise, but except to the extent this impacts on his epistemology that is not a subject I'll discuss here. In chapter 9 I'll discuss Levi's decision theory, but here I'll focus more narrowly on his epistemological innovations.

7.1.1 *Levi's Argument for Imprecision*

The argument Levi (1974)¹ gives for allowing imprecision is as follows:

- (1) Replacement in bodies of beliefs is only rational if it can be construed as a contraction followed by an expansion.
- (2) It is sometimes permissible to replace credal probability functions.

Therefore,

¹ Unless otherwise stated, all references in this chapter are to the subject of the section in which the reference occurs.

- (3) Sometimes it is permissible to not rule out more than one credal probability function. That is, vagueness must be allowed, at least between the times that a rational agent has one precise credal probability function and the time she has a different one.

Although I agree with the conclusion, I disagree with each of the premises. Since my objections to (1) will show that the argument is unsound, I won't say much about (2). If the idea behind (2) is that evidence can show us our choice of probability function was in error, it is clearly a mistake. One of the conditions on a function being reasonable is that its conditionalisation under any evidence at all is reasonable. Levi occasionally argues something like this, but in later works, particularly Levi (1980), the motivation for (2) is a pragmatist theory of belief. Different probability functions, he says, might be warranted if our values change. Without a wholesale discussion of what's wrong with pragmatist epistemologies, I couldn't give a reasonable account of my objection to (2). So I'll stick to discussing (1).

The motivation for (1) comes from Levi's non-probabilist epistemology, so for a while I can simply discuss that. In Levi's theory, rational agents have belief sets which are sets of sentences closed under entailment. So rational agents believe all the logical truths, as well as all of maths and set theory. They also believe that each one of their beliefs is an item of knowledge. For this reason Levi refers to the sets as bodies of knowledge. This is a little misleading since many of an agent's beliefs are clearly not knowledge for the simple reason that they are false. Given this use of words, we can't take Levi's claim that knowledge is infallible at face value. What he means is that if an agent takes herself to know p , in effect believes p , then $\neg p$ is not, for her, possibly true (1980: 13), where 'possibly' is interpreted epistemically. In more familiar terms, anything which contradicts our beliefs is not a doxastic possibility.

Now this might seem to conflict with some venerable pragmatist doctrines, particularly the doctrine that all of our beliefs are open to revision. In his words, there is a worry that Levi's

infallibilism implies incorrigibilism. That, he assures us, is not the case, though as mentioned he does in effect take beliefs about the ‘conceptual scheme’, i.e. maths and set theory, to be incorrigible. He allows that corpora can be contracted, and hence allows that agents can give up beliefs. But there is a worry as to why an agent would give up a belief which was, to her mind, necessarily true. One reason is that the agent might have ended up with an inconsistent corpus of beliefs, say by coming to believe some observation sentences. This is clearly plausible, but it’s not obvious that any other good reason exists. Certainly, Levi’s attempt to motivate other grounds seems quite weak.

Other good reasons exist for contracting a corpus. Suppose the initial corpus contains some theory T_1 . A second theory T_2 contains T_1 . From X ’s initial point of view, T_2 is certainly false. Yet it may be superior in all other respects to T_1 as a means for furnishing systematic explanations in some domain ... In such cases, X might be prepared to suffer a loss of information due to the removal of T_1 from his corpus in order to be in a position to take the truth of T_2 to be seriously possible.

(1980: 60)

However, persuasive this sounds in the abstract, in real cases it sounds implausible. Remember that T_2 isn’t just a theory the agent doesn’t believe, it is a theory believed to be false. Now try to imagine a circumstance in which the elegance or systematicity of a theory you believe false, say Aristotle on any physical science, or perhaps Marx on any social science, could make you give up beliefs you currently have. Maybe I’m just dogmatic, but to my mind if a theory is known to be false it doesn’t matter how pretty it is. Still, the point about inconsistency in the corpus necessitating contraction probably is enough here. Levi notes that inconsistency can easily arise if we make rule-governed expansions, so we need contraction for this reason.

He also allows that beliefs can be replaced. An agent can replace a belief in p with a belief in $\neg p$. “Replacements are shifts from corpora to other corpora inconsistent with the initial ones.” (1980: 63) The history of science, as well as everyday life, tells us that replacements often happen. But their justification is, if anything, even weaker than the justification of contractions. From the agent’s point of view, replacement amounts to adopting something false as a belief, not just giving up a true belief. Levi claims we can get a justification if we ‘decompose’ (his term) the replacement into a contraction followed by an expansion.

We might interpret this in two ways. First, Levi might be making the empirical claim that all reasonable replacements consist of contractions followed (at a later time) by an expansion. So the decomposition might be a closer analysis of what actually happens. Alternatively, he might be claiming that all reasonable replacements can be rationally reconstructed as justifiable contractions followed by justifiable expansions, even if these aren’t temporally distinct in the mind of the agent. The text isn’t particularly clear on which he intends, but I think each possibility is worth investigating. I claim the first turns out to be false, and the second turns out to be unjustifiable.

We can see that the empirical decomposition claim is false by looking at everyday examples. Jack reads in the morning newspapers that Smith will be the starting pitcher for the Yankees in tonight’s game. He hears this repeated on the lunchtime news. Jack is a fan of Smith, so he decides to go to the game to watch Smith. As it will turn out that Smith isn’t the starting pitcher, we can’t say Jack knows Smith will be the starting pitcher, but it is something we should say is in his corpora of beliefs as Levi puts it. The evidence for this is that Jack believes he knows this, he will answer the question ‘Who will start pitching for the Yankees?’ with the answer ‘Smith’ and he acts as if he knew Smith was the starter².

² It might be objected by some that the criteria for belief being used here is too weak. Even if this were true,

Unbeknownst to Jack, Smith gets injured in the pre-game warmup and Jones becomes the starting pitcher. Jack only realises this when he walks into the ground, hears on the ground public address system that the opening pitch is about to be thrown, looks to the mound and sees Jones pitching. He quickly comes to believe, on the basis of his overwhelming sensory evidence, that Jones is the night's starting pitcher. Does he first drop his belief that Smith will start and then adopt the belief that Jones starts? No; the two happen at exactly the same time. And this is quite reasonable. So Levi's construal of Jack's epistemic dynamics can be at best a rational reconstruction.

If, however, we take the rational reconstruction route, the purported justification of replacement becomes implausible. Levi notes that when an agent removes p from their corpus, they leave open the possibility that they will at a later time expand to include $\neg p$, which they believe to be false. Hence every contraction creates the possibility of expanding into (perceived) error. In replacement this possibility is realised. Against the claim that this shows all contraction to be unjustified, Levi claims that agents are allowed to be 'myopic' (again his term) when contracting. Agents are allowed to simply ignore the long-run effects of their actions, in this case the possible error to which they'll be led. The defence of this is simply that the alternatives are worse. "I cannot prove that I am right in singing the praises of myopia. Nevertheless, the alternatives seem far less attractive." (1980: 71)

The problem with the myopia account is that it is completely implausible on the rational reconstruction model. The idea behind it is that when the agent expands her beliefs to include $\neg p$, she no longer believes p , so she isn't coming to believe a falsehood. And she is allowed to

it's not something Levi can claim, as he allows beliefs on much weaker evidence than this. In one example he allows an agent to believe p even though she knows the objective chance of $\neg p$ is positive (1980: 275). In another he lets the agent infer p although her credal probability for $\neg p$ is 0.09 (1980: 136). The consistency of this approach seems dubious to me, but I would need much more space to address all the possible arguments here.

myopically ignore the possibility that she will do this when she originally contracts. But on the rational reconstruction view, her beliefs do include p at the very time (or at least right until) she believes $\neg p$. Given this, she has to myopically ignore a possibility that's actually happening at the time she contracts. This is implausible, so I conclude something's wrong with Levi's account of replacement. And this implies that this argument for (1) fails, so he has no reason to infer (3).

7.1.2 *Levi's Calculus*

Despite all that, (3) is correct, so it is worthwhile looking at how Levi incorporates this into his account of rational belief. For Levi, rational belief states can be represented by sets of probability functions. Anything the agent believes receives probability 1 according to each of these functions, though the converse need not hold. The functions are two-place, so conditionalisations on propositions with zero credal probability are defined, but not conditionalisations on propositions inconsistent with what the agent believes (1980: 221). The functions need only satisfy finite additivity; Levi claims the arguments for countable additivity in the literature are generally question-begging (1980: 224-7).

His attitude to updating is a little confusing at first glance. "I myself am willing to endorse *confirmational conditionalisation* even though I reject *confirmational tenacity*. Consequently, I do not think that shifts in credal state due to expansion should be *temporal credal conditionalisations* in all cases." (1980: 85, my italics). In simple terms, he adopts a position similar to van Fraassen (1989), where agents are not required to have exhaustive rules for updating, but to the extent they have rules they shouldn't conflict with conditionalisation. However, the story is a bit more complicated than that.

Levi does think that (ideally rational) agents should have 'commitments' as to how to update under any information, and these should obey conditionalisation. This is his rule of confirmational conditionalisation. However, these commitments are, like everything else, subject to revision. The

agent need not tenaciously hang on to them in all circumstances. To insist on this would be to adopt confirmational tenacity, and Levi rejects it. So an agent can have a ‘commitment’, but not be committed to implementing it³. This strikes me as an unusual use of *commitment*, and my use of scare quotes is intended to be a reminder of this. Hence when some evidence comes in the agent can, at that stage, choose to drop some commitments, and hence the new epistemic state that’s adopted need not be the conditionalisation of the old one on the evidence. This is what he means when he denies all shifts are temporal credal conditionalisations. If we restrict the term *commitment* to those ‘commitments’ the agent is committed to implementing, and insist the agent can’t both have a ‘commitment’, and be committed to *not* implementing it, we get van Fraassen’s position as I set out in the previous paragraph.

In chapter 5 I allowed that any precisification of a reasonable epistemic state was itself reasonable. Levi appears to adopt a different position. He explicitly endorses “the contention that one should not rule out [probability]-functions unless one has a warrant for doing so.” (1980: 89). However, he allows that warrant may be interpreted liberally to include the agent’s values and goals. So the difference here is not as wide as it first appears. My argument for allowing arbitrary precisification turns on my differences with Levi’s decision theory, which I’ll discuss in the next chapter.

The most interesting aspect of Levi’s calculus is that he insists reasonable epistemic states be represented by *convex* sets of probability functions. So the set must be closed under linear mixtures provided the weights are all non-negative. There is a *prima facie* argument against this. As noted in 3.6.6, Jeffrey (1987) argues that the epistemic state containing just the information that *A* and *B* are probabilistically independent is reasonable despite not being convex.

³ It isn’t entirely clear from Levi’s text, but possibly we could read this as saying agents must be disposed to update by conditionalisation, but this disposition may be finkish.

Prima facie claims, however, do not settle the matter, so I ought to look at Levi's arguments for convexity. His arguments rely on his decision theory, which I think is mistaken, but I'll ignore that complication here. The problem is that the argument he gives doesn't go far enough. He argues that the epistemic state represented by just the two probability functions, $Pr_1(p) = 0.4$ and $Pr_2(p) = 0.6$ (we are only interested here in the truth or otherwise of p), is unreasonable. I agree; indeed, given the Equivalence Analysis, I am at a loss to know what we might mean by saying this is an agent's epistemic state.

The problem for Levi is that even if states like this one are unreasonable (or meaningless) this doesn't clinch the case for convexity. And he openly admits that he has 'no proof' (1980: 192) for requiring convexity beyond cases like this one. Consider the following property of sets of probability functions.

Continuity: A set S of probability functions is continuous iff for any proposition A and any numbers x, y, z , if there exist $Pr_1, Pr_2 \in S$ such that $Pr_1(A) = x$ and $Pr_2(A) = y$ and $x \leq z \leq y$, then there exists a $Pr_3 \in S$ such that $Pr_3(A) = z$.

A set of probability functions satisfies Continuity (for short, is continuous) iff the values of $Pr(A)$ takes for Pr in that set is an interval for every proposition A . The set which Levi argues is unreasonable isn't continuous. However, the set Jeffrey argues is reasonable is continuous. This suggests that the problem with Levi's set isn't its lack of convexity, but its lack of continuity. I think the meaning considerations from chapter 3 make continuity a plausible constraint, but its purpose here is simply to show that Levi hasn't refuted Jeffrey's *prima facie* objection to convexity.

7.2 van Fraassen

The most obvious debt this dissertation owes to Bas van Fraassen's work is his development and promotion of supervaluations. Not only did he do much to develop the technique and promote it

within the philosophical community, (particularly in van Fraassen (1966)), he has also argued that it ought be applied in the case of imprecise credences (van Fraassen (1990)). Some of the most interesting applications of this idea to date are also due to van Fraassen. In van Fraassen (1989) he uses this idea to develop an analysis of agnosticism to make his agnosticism about unobserved entities plausible, and in van Fraassen (1995) he uses the possibility of vagueness to fend off some objections to his principle of Reflection. Here I want to concentrate on an interesting technical result which I flagged in chapter 3, and which might have some technological implications.

7.2.1 *Why Imprecision Is Allowed*

In this dissertation I've taken the line that even perfectly rational agents may have imprecise degrees of belief. van Fraassen isn't as much concerned with what perfectly rational agents think as with what everyday agents think. If we had precise credences, he notes, we could only specify our epistemic states with an infinite number of judgements. However, since we are finite beings, "our expressible opinion must be expressible in a finite number of judgements" (1990: 353). Hence our degrees of belief must be imprecise.

It's not altogether obvious that this is correct. Provided that we take a dispositional view of opinions, and we individuate dispositions finely, it seems we might have an infinite range of opinions. The only requirement is that a single physical state must be capable of instantiating multiple dispositions. This certainly looks possible. Given van Fraassen's interests, though, we don't need this argument. For precision requires that our opinion include non-trivial judgements about every subject imaginable, and this is clearly false. The possibility I referred to is simply irrelevant because it isn't instantiated. So someone who is interested in how humans with their limited resources and interests should reason, or more generally an investigator into epistemological norms who places some weight on the dictum 'ought implies can', should take vagueness in credences seriously.

7.2.2 *van Fraassen's Calculus*

What follows is a brief summary of the proof in (1990) that what van Fraassen calls figurehood is preserved under conditionalisation, and some comments on the potential importance of this proof. A random variable g is a function from a possibility space to reals. Given a probability function Pr , we can work out the expected value of g according to Pr , as $\sum Pr(w)g(w)$, where g ranges across the possibility space. This will be notated as $E_{Pr}(g)$. When the possibility space is infinite we need to be more careful to ensure this sum can be calculated. van Fraassen makes all the necessary assumptions, but for this sketch I'll sacrifice a little precision for brevity. Readers who know the difficulty will also know how to make the corrections⁴. We can add random variables to constants or to each others, and we can multiply them by constants (or indeed each other, though this is less important). In each case the addition or multiplication is simply done possibility by possibility.

Using expected values we can give a standard form for judgements. Say $E(g \geq a)$ is satisfied by Pr iff $E_{Pr}(g) \geq a$. Then a judgement is a Boolean combination of statements of the form $E(g \geq a)$. From this we can get $E(g \leq a)$, which means the same as $E(-g \geq -a)$, and using these two we can express many other judgements. We can interpret ordinary judgements (like 'A is more probable than not') using indicator functions. The indicator function of A is a random variable I_A which takes value 1 if w is in A and 0 otherwise. We get the following results:

$$Pr(A) = r \text{ iff } E_{Pr}(I_A) = r;$$

$$Pr(A) : Pr(B) = r \text{ iff } E_{Pr}(I_A - rI_B) = 0$$

⁴ There is one aspect in which van Fraassen is considerably more precise than most of the literature. The basic possibilities include information about what we are believing. This necessitates quite a few complexities which are often ignored; it is more usual (but less accurate) to write as if each element of the possibility space is independent of our opinions.

$$Pr(A | B) = r \text{ iff } E_{Pr}(I_{A \cap B} - rI_B) = 0$$

Using these results van Fraassen shows how to express many probabilistic judgements. For example we have the following:

$$\text{'}A \text{ is at least as likely as } \neg A\text{'}: E(I_A - I_{\neg A} \geq 0)$$

$$\text{'}A \text{ is } r \text{ times as likely as } B\text{'}: E(I_A - rI_B = 0)$$

$$\text{'Given } C, A \text{ seems twice as likely as not'}: E((I_{A \cap C} - (2/3)I_C) = 0)$$

There are, however, some judgements we can't express. For any function g , the set of probability functions satisfying $E(g \geq a)$ is convex. To prove this simply note that the value of g according to an equal mix between two functions is simply the average of its value according to those functions. And since the intersection of convex sets is itself convex, we can't express epistemic states represented by non-convex sets of probability functions. So for the reasons given in the previous section, this language looks like it can't be expressive enough to say all we want about partial beliefs. Nevertheless, it has interesting properties.

If we take the standard form judgements as basic, we can define an agent's epistemic state as the set of probability functions which satisfy these judgements. A *figure* is defined as either a finite set of judgements of the form $E(f \geq 0)$, or equivalently the set of probability functions which satisfy all these judgements⁵. The complexity of a figure is the smallest number of judgements of which it is the intersection. That our mental states are finite is reflected in the claim that (coherent) human epistemic states are figures.

We now get a pair of interesting questions. Is the property of figurehood preserved under conditionalisation, and what effect does conditionalisation have on complexity? van Fraassen

⁵ Since $E(f \geq a)$ is the same as $E(f - a \geq 0)$ there is no loss of generality from this definition.

provides a nearly constructive proof that the complexity increases by at most two. First some notation. For any set \mathcal{P} of probability functions, let \mathcal{P}_B be defined as $\{Pr(\cdot | B): Pr \in \mathcal{P} \text{ and } Pr(B) > 0\}$. The second conjunct in this definition is a little restrictive, but this assumption is used in the proof. Let $|B|$ be $\{Pr: Pr(B) = 1\}$. We will now construct \mathcal{P}_B . Note that it follows immediately from the definition that $(\mathcal{P} \cap \mathcal{Q})_B = \mathcal{P}_B \cap \mathcal{Q}_B$, so we can look at single half-spaces. That is, we'll assume $\mathcal{P} = E(g \geq 0)$.

Every Pr has an orthogonal decomposition in terms of B for $0 < Pr(B) < 1$:

$$Pr = cP^+ + (1 - c)P^-, \text{ where } 0 < c \leq 1, P^+ = Pr(\cdot | B), P^- = Pr(\cdot | \neg B).$$

It follows that $E_{Pr}(g) = c E_{P^+}(g) + (1 - c) E_{P^-}(g)$. If P^+ exists then $Pr(B) > 0$. So $q \in \mathcal{P}_B$ iff $q = P^+$ for some $Pr \in \mathcal{P}$. And $Pr \in \mathcal{P}$ iff $E_{Pr}(g) \geq 0$, i.e. $c E_{P^+}(g) + (1 - c) E_{P^-}(g) \geq 0$. This can be trivially rewritten as $c \cdot E_{P^+}(g) \geq -(1 - c) E_{P^-}(g)$.

Now there are two cases to consider. Case 1, there is a probability function $q \in |\neg B|$ such that $E_q(g) > 0$. Then for any $p \in |B|$ there will be a value of c in $(0, 1]$ such that $c E_p(g) \geq -(1 - c) E_q(g)$, since when c goes to 0 the LHS will go to 0 and the RHS to $-E_q(g)$. Let c be such a value. Hence $Pr = cp + (1 - c)q$ is an element of \mathcal{P} since p and q are its decompositions. Since $Pr(\cdot | B) = p$, and p was an arbitrary member of $|B|$ it follows that $\mathcal{P}_B \subseteq |B|$. But since every function in \mathcal{P}_B must satisfy $Pr(B) = 1$, $\mathcal{P}_B \supseteq |B|$, so $\mathcal{P}_B = |B|$.

Case 2 then is when there is no $q \in |\neg B|$ such that $E_q(g) > 0$. In that case for $p \in |B|$ (and any $q \in |\neg B|$), there will be a value of c in $(0, 1]$ such that $c E_p(g) \geq -(1 - c) E_q(g)$ iff $E_p(g) > 0$. If we do have $E_p(g) > 0$ then again $Pr = cp + (1 - c)q$ is an element of \mathcal{P} since p and q are its decompositions. However, if $E_p(g) < 0$, then for any c, q , if $Pr = cp + (1 - c)q$, $E_{Pr}(g) = c E_p(g) + (1 - c) E_q(g) < 0$, so $Pr \notin \mathcal{P}$. This implies $\mathcal{P}_B = |B| \cap E(g \geq 0)$.

So for any half space $E(g \geq 0)$, conditionalisation on B takes it either into $|B|$ or its intersection with $|B|$. The former possibility will occur if there is an element q of $|\neg B|$ such that $E_q(g) > 0$, the latter otherwise. In the general case, when P is a set of half spaces, the same story holds. P_B is the intersection of those judgements $E(g \geq 0)$ such that there is no element q of $|\neg B|$ such that $E_q(g) > 0$, with $|B|$. And we can write $|B|$ as the intersection of $E(I_B - 1 \geq 0)$ with $E(1 - I_B \geq 0)$. Hence the complexity of P_B is at most two more than the complexity of P .

What is most interesting about this proof is the possibility it opens up for computational purposes. If there is a simple test to determine which case applies for a given judgement, the one non-constructive part of van Fraassen's proof, then it will be easy to conditionalise vague epistemic states that are figures. Now we noted that the language of figures isn't quite as expressive as we might like, but given how easily they can be updated, this loss might not be excessive.

7.3 Jeffrey

In Jeffrey (1983a) there are two reasons given for allowing degrees of belief to be vague, or as Jeffrey puts it, for allowing reasonable epistemic states to be representable by "propositions", which are just sets of probability functions. The first is technical; even if we assume an agent has a complete preference ordering over all possible states of affairs, this may be consistent with an infinite number of pairs of subjective probabilities and utilities. (See Jeffrey (1983b) for the details.) The second is more pragmatic; real people simply don't have enough preferences or dispositions to make it the case that their epistemic state is representable by a single probability function.

Part of the attraction for Jeffrey of this position is that it allows for more tools in the Bayesian toolkit than just conditionalisation⁶. For example, let A_i , for $1 \leq i \leq n + 1$ mean the i th trial was a

⁶ One of the motivations behind Jeffrey's paper is to respond to Clark Glymour's definition of Bayesians as those who believe all updating should take place by conditionalisation.

success, and A_i be the indicator function of A_i . We needn't bother what the trial is, it might be tossing a coin and seeing if it lands heads, or dropping a plate and seeing if it shatters. Assume that we know there have been s successes on the first n trials, and for whatever reason think the probability of success on a given trial is constant.⁷ Hence our epistemic state will be representable by the set $\{Pr: Pr(A_i) = Pr(A_j), 1 \leq i \leq j \leq n+1\}$. Since we know that $A_1 + \dots + A_n = s$, our expectation for $A_1 + \dots + A_n$ ought be s . But that expectation is just $Pr(A_1) + \dots + Pr(A_n)$; as each term of that expression is equal and they must sum to s , each term must equal s/n . And since $Pr(A_{n+1}) = Pr(A_1)$ for all Pr in our representor, $Pr(A_{n+1}) = s/n$. This, thinks Jeffrey, is an important part of the explanation of the importance of frequency, and it's all been done without conditionalisation, and indeed without saying anything about the probability of success on one trial conditional on the success or failure of another.

The results Jeffrey gives in this area, as can be seen from the above, are hardly of stunning technical importance, and his justification is not particularly interesting from the point of view of this dissertation, as I'm interested in arguments to the conclusion that even ideally rational agents should have imprecise degrees of belief. However, Jeffrey's endorsement of vagueness is historically important for two (related) reasons. First, Jeffrey is arguably a paradigm case Bayesian; hence, it is inappropriate for supporters of imprecision in probability to say that the enemy are the Bayesians, as is occasionally done. Secondly, it seems Jeffrey's paper has been more influential in converting philosophers to the view that imprecision is rationally permissible than any other. Hence it is worthwhile noting here.

⁷ In the kind of examples that Jeffrey seems to be considering, "textbook trials", this hardly seems like a reasonable step. To see why, consider what happens to the reasoning in the text when s equals n and is rather low.

7.4 Kyburg

In many ways the theory of probability developed by Henry Kyburg over a number of years is the most similar to the one defended here. In particular he argues that probability is an objective relation between a claim and the evidence for it and that the value of this relation is non-contingent, legislative for rational belief, and not necessarily numerical. However, the differences between his position and mine are also substantial. His is a logical, not merely a necessitarian, theory of probability. On Kyburg's story, probability, like provability, is a meta-linguistic relation which holds between sentences (not propositions) by virtue of their syntactic properties. As a consequence of this, probability relations are language dependent. Further, the interval-values that probability relations take are in a sense primitive or unstructured; there's nothing to say about the probability of p , beyond the fact that it is the interval $[x, y]$. In the theory defended here, it matters (particularly for comparatives) which probability functions generated which values in $[x, y]$. Most importantly, all knowledge of probabilities is grounded in knowledge of frequencies. This is not the same as analysing probability as frequency, but it does have rather similar consequences. I have been deliberately agnostic on the question of how we come to know probabilities; my main aim in this section will be to defend my agnosticism against Kyburg's particular religion.

This is Kyburg's summary of his position, which has remained unchanged on these questions for over 30 years.

Roughly speaking, we shall say that a statement has the probability p (in general p will be an interval i of reals), relative to a body of knowledge, when (a) it is known in that body of knowledge that the statement is equivalent to (has the same truth value as) a statement of the form $\lceil a \in b \rceil$; (b) it is known in that body of knowledge that a belongs to c ; (c) that body of knowledge contains the statistical knowledge that the proportion of objects in c that belong to b falls in

the interval i ; and (d) there is nothing in our body of knowledge that conflicts with this assignment of probability. (Kyburg 1974: 156-7).

The technical construal of (d) has changed in response to some problems the early account generated, and I'll outline its evolution below. As should be clear from this quote, Kyburg's construal of evidence is much different to mine. For one thing, he takes evidence to be sentential, not propositional. For another, it includes 'statistical' sentences, such as sentences about the proportions of heads among the tosses of fair coins. Kyburg takes a fundamentally different position to the one (implicitly) adopted here on the nature of induction. I construe inductive arguments as concluding that something is probable. Kyburg takes them as having a categorical conclusion, but with reasoning that is defeasible, or perhaps merely probable. (The distinction between these two accounts of induction is given most clearly in Hempel (1965).) Hence Kyburg thinks we can include the conclusions of our inductions in our evidence set. These inductive conclusions generate statistical statements about frequencies, and these generate probability statements.

This leads Kyburg to make a more dramatic departure from conventional wisdom about rational belief sets. Kyburg thinks that we should take as beliefs anything that we believe to at least a certain, high, degree, say n . Now imagine a (fair) lottery with more than $1/(1 - n)$ tickets. For every ticket, we will believe to greater than degree n that it will lose, hence we believe it will lose. However, we also believe that some ticket will win. Hence we are committed to believing a set of inconsistent propositions. This is the lottery paradox, itself due to Kyburg (1961). Many writers have taken this as an argument against construing (full) belief as high degree of belief. Kyburg takes it to be an argument against requiring rational belief sets to be logically closed. Many of the technical complications in his theory arise from this. For example condition (a) above can't be construed as saying $p \equiv (a \in b)$ is in our evidence set. Rather we construe it as saying there is some chain of biconditionals in our evidence set that starts with p and ends with $a \in b$. And in

fact we have to say the same thing about $a \in b$. It might be that our evidence set merely contains $a \in d$ and $d \subset b$; this has to be sufficient. The complications induced this way don't seem formally necessary to Kyburg's theory (when explaining some technical aspect of his theory he will occasionally assume logical closure for ease of exposition) but they are I think philosophically necessary. If we agree that we can have full belief in statistical statements (which he requires) then we are forced into his resolution of the lottery paradox and hence must deny closure.

As I said above, the main technical move in his theory has been the changes in condition (d), the randomness condition. Since Kyburg takes an epistemic construal of probability, he seems justified in interpreting randomness as ignorance. The question is just what type of ignorance it is. The probability of $a \in b$ is the interval i iff it is known that a is a random member of the class c , and i is the smallest interval such that the proportion of c 's in b is known to be in i . If all these conditions hold, we'll call c the reference class for the probability judgement. As a first approximation, we can read the underlined part of the clause as saying that we don't know anything special about a other than its membership of c . The problem is now to say what is 'special' knowledge. Again as an approximation, we know something special about a if we know a is a d , and we know the proportion of d 's in b is j , and i is not a sub-interval of j . This, Kyburg hoped in (1974), provided the best trade-off between using the smallest reference class available, and using the most precise knowledge we have.

To illustrate, assume we know of a_1 that it is a d , and that $d \subset c$. We know the proportion of d 's in b is $[0.3, 0.8]$, and the proportion of c 's in b is $[0.7, 0.9]$. Since neither interval is included in the other, the information we get about the probability of a_1 being a b from examining the d 's and examining the c 's clashes, and in this case we use the smaller reference class. Hence the probability, on our evidence, of a_1 being a b is $[0.3, 0.8]$. Compare this case with a_2 , which is known to be an e , and it is known that $e \subset c$. The proportion of e 's in b is known to be in $[0.6, 1]$. In this case the information we get about a_2 's probability of being a b from looking at the e 's and

looking at the c 's does not clash; one is merely more precise than the other. In this case we don't worry about the size of the reference classes, we just use the more precise knowledge. So the probability of a_2 being a b , according to our evidence, is $[0.7, 0.9]$.

There are more complications (dealing for example with a_3 which we know is a d and an e), but that's the basis of the randomness criteria in (1974). However, it turns out to lead to rather odd results. With the benefit of hindsight, these oddities seem inevitable given the two-part rule adopted. Some cases are always going to fall through the cracks. Let c be the set of all draws made from an urn of counters, and d the set of all draws made yesterday. (Hence $d \subset c$.) Let a be a counter that was drawn yesterday. Our knowledge of the frequency of colours drawn amongst c and d is given in this table.

	White	Red	Blue
c	$[0.5, 0.5]$	$[0.25, 0.25]$	$[0.25, 0.25]$
d	$[0.4, 0.4]$	$[0, 0.6]$	$[0, 0.6]$

When we are determining the probability that a is white, the information about c and d clashes, so we use the smaller reference class. Hence the probability that a is white is precisely 0.4. For red (and *mutatis mutandis* for blue) the information about c and d does not clash, so we use the reference class that gives us the more precise information. Hence the probability that a is red (blue) is 0.25. Kyburg is committed to saying that whenever probabilities are precise they should obey the probability calculus (these are just set-theoretic tautologies when interpreted as statements about frequencies), but here he is committed to a breach of that calculus. The solution (in Kyburg (1983)) is to change somewhat the definition of randomness. We now use a smaller reference class if it is known to generate clashing information about the predicate in which we are interested, or any other (known to be) disjoint predicate. Getting this formally right requires quite a lot of technical work – in particular now we have to watch out that this rule isn't only obeyed

for the partition of the possibility space we are using, but also for any refinement of it – but this captures the philosophically important aspects of the amendment.

There has been quite a lot written about the technical aspects of Kyburg's work, and I don't particularly want to add to it here. Rather, I want to question the philosophical presumptions of his theory. In particular, I don't think the reliance on frequencies and the language-dependence of Kyburg's theory can be defended when we look at how probability is (and ought be) applied to the actual world.

When we are trying to discover the probability of interesting claims about the real world (as opposed to the colour of a marble in an urn) Kyburg notes that we may need to be creative in discovering the relevant frequencies (1961: 266ff). For example, what makes it true that 'It is probable (given our evidence) that Caesar crossed the Rubicon' is not the known frequency of Caesar-type leaders crossing Rubicon-type rivers. Rather it is, for us, the frequency that historical assertions assented to by the vast majority of historians turn out true. That frequency we (apparently) know to be high, and hence we can give a high probability to 'Caesar crossed the Rubicon'. For the historians themselves, who could hardly reason this way, they know there are records from the time which say the Caesar crossed the Rubicon, and they know the frequency of such records being correct, and this frequency underlies their probability judgement.

However, even with such ingenuity, I doubt it will be possible to discover enough frequencies to ground all the seemingly justified probability judgements that are made. For non-experts, the move of referring to the frequency of experts being correct can always be used⁸, so I'll look at how experts should reason. Consider one rather recent case. In early 1998 President Clinton was entangled in a scandal over an affair he'd apparently had with a young White House aide, Monica

⁸ Or at least attempted; in many fields I remain highly skeptical that experts are in general correct, particularly in prediction.

Lewinsky. This was of itself a minor scandal, but the real danger was that according to some allegations, he had lied about the affair under oath, and encouraged Ms. Lewinsky involved to do the same. If this were true, he'd have both committed and suborned perjury, and many felt this was grounds for impeachment. Given all this, Washington correspondents were frequently asked for their assessment of the chances of Clinton being impeached, or perhaps for his relative chances for survival on different strategies.⁹

I claim that Kyburg cannot explain the use of probability in these reports from Washington, and that these reports are, if a little charitably interpreted, correct uses of 'probability'. This is a problem for a theory such as Kyburg's which does aim to capture uses of probability in natural language. More generally, I think the probabilities being discussed were legislative for rational belief, but were not based on frequencies, which if true would remove a foundation stone of Kyburg's theory.

As the crisis unfolded, one of the internet-based news magazines¹⁰ listed its daily assessment of the chances of Clinton being impeached. I interpret these as claims as to the probability of Clinton being impeached given all the evidence publically available¹¹. Their assessment of this probability, over the eleven main days of the crisis, was as follows.¹²

⁹ At the time of writing, July 1998, the White House seems to have successfully swayed public opinion by focusing on procedural matters in the investigation. Hence there is little political will for impeachment outside the usual Republican suspects. But what I said in the last footnote holds for my predictions too, so don't hold me too tightly to that guess.

¹⁰ *The Slate*, at <http://www.slate.com>.

¹¹ They were explicitly estimates of the 'chance' of impeachment, which amounts to the same thing. They weren't particularly timely measures of chances because of the lag between something happening which affects the chance and it becoming publically known, but this doesn't affect the comments in the text.

¹² As background, between the 21st and 25th more rumours about evidence of the affair appeared, and erstwhile Clinton supporters publicly discussed impeachment and, occasionally, urged a resignation. From

Day (of January)	21	22	23	24	25	26	27	28	29	30	31
Probability of Impeachment	25%	35%	40%	45%	45%	40%	38%	35%	34%	30%	21%

Now the real numbers are obviously fuzzier than this, and the particular choice was hardly made seriously. In particular we should agree with Kyburg that the probability of impeachment on a given day should be an interval rather than a number, and perhaps quite a broad interval at that. What I do think should be taken seriously are the comparatives. Except perhaps for the 1% change from the 28th to the 29th, all the other day-to-day comparative statements look correct to me, given the evidence that was available at the time. I don't extend this to comparisons between the various probabilities quoted on the way up and the way down.¹³ So it was more probable on the 23rd he'd be impeached than the 22nd, equally (or perhaps incomparably) probable on the 25th that he'd be impeached as on the 24th, and just about every day after that it was getting less probable.

I think a successful theory of probability should be able to explain these comparative statements, and be able to say something about the absolute probability statements. I doubt Kyburg's theory could explain the comparatives, but I must be a little cautious because Kyburg has not, as far as I can tell, written anything on comparative probability sentences. The strongest theory of

the 26th onwards Clinton was supported as more and more of the rumours turned out to be false, and polls showed that not only were most people against impeachment at present, a majority opposed impeachment even if the charges were proven, and at the end of the month Clinton's approval rating hit an all time high.

¹³ It's easy with hindsight, but I don't know how Clinton's position could have been considered worse on the 27th than the 22nd. Once things started looking good for the President, the media, in order to protect a great story and justify their interest in it, overestimated the problems Clinton was in. Hence I'll only use day-to-day comparisons, because they seem the only justified ones.

comparatives he could plausibly endorse is the following. Say the probability of p is $[x_p, y_p]$ and the probability of q is $[x_q, y_q]$. Then the probability of p is greater than the probability of q iff $x_p > x_q$ and $y_p > y_q$. (Perhaps he could say that one of these is allowed to be an equality, but that doesn't sound overly plausible.) Given the way his theory is developed, it would seem more appropriate to say that the probability of p is greater than the probability of q iff $x_p > y_q$, or $x_p = y_q > x_q$.

Since what I say is least effective the stronger the theory of comparatives Kyburg adopts, I'll assume he adopts the stronger of the two theories mentioned above. Now Kyburg has to find no fewer than nine different reference classes for judging the probability of 'Clinton will be impeached' over the eleven days of the crisis. If the reference class used on day t is the same as that used on day $t + 1$, then he can't say the probability of impeachment on day t is different (either greater than or less than) the probability of impeachment on day $t + 1$. I doubt that he could reasonably find one reference class; even if I'm wrong about this, I'd have to be wrong in a big way for him to find enough reference classes to make every comparative turn out true.

What reference classes could be used? The class of Presidents who are impeached and removed from office doesn't help; it's still empty. The class of events predicted by Washington correspondents might help for explaining the probability for some people, but not surely for the Washington correspondents we are looking at. If we knew that impeachment was an element of the class 'political actions desired by a majority' we might get somewhere, but that was never known, and as it turned out never true. (It was the realisation of this that sent the probability of impeachment diving at the end of January.) Perhaps instead of just looking at the frequency that Presidents are impeached, we could look at impeachment rates for elected officials (e.g. state governors). But it is hardly plausible to say the potential impeachment of a President is a random member of the class of potential impeachments of elected officials.

There is a more general difficulty for Kyburg here. Even if we had a long string of impeached Presidents, so we could judge the frequency of impeachment once a scandal got to such and such a stage, it's hard to see how that could be relevant here. One of the distinguishing features of this crisis was the massive involvement of the media, particularly internet-based media. And by internet-based media I include traditional newspapers who were concentrating on getting stories onto their web sites as quickly as possible, not just the best story for the morning newspaper. Some of the most absurd rumours of the scandal appeared through these sources. But these rumours not only affected everyone's perception of the probability of impeachment, there was every chance that, by affecting the way the crisis was viewed, they could have had a real effect. This affair was nothing like anything that had preceded it, but that didn't prevent it being possible to make reasonable probability judgements. Indeed, anyone who had all the evidence but didn't form judgements about the probability of impeachment of similar form to that given above would be unreasonable. By similar form I just mean believing impeachment was more probable at the peak of the crisis, around the 25th, than it was a few days before and after, and believing that at these times impeachment was more probable than it was one month previously. But if Kyburg's correct, we oughtn't to be able to make probability judgements about such *sui generis* events.

The theoretical difficulty for Kyburg this raises is that it seems we may have a use, in social sciences at least, for 'pathological' predicates. For his theory to generate decent results, Kyburg needs to rule out pathological predicates like grue. This is related to both the language-dependence of his theory, and the need to generate known statistical statements as building blocks. However, in social sciences, and generally in making probability judgements about everyday life predicates that are broadly pathological may be needed. I don't mean we necessarily need predicates like 'grue' which change, in some sense, their meaning at the turn of a millennia.¹⁴ What we do need

¹⁴ Though we should never underestimate the importance of millennial symbolism in politics. Australia is currently debating whether it should change from a monarchy to republic on the first day of the new

are predicates which, in some sense, shift a little when real factors in the outside world shift. From the political scientist's perspective, President Clinton's having an affair in this age of mass media is a quite kind of different occurrence from, say, President Kennedy's having an affair. Predicates which behave like 'green before an internet exists, blue otherwise' have some importance. Kyburg claims (1990: 126-130) that his theory can deliver the tools to develop the language of science, and this language will help us avoid the grue pitfalls. For the social scientist this seems both impossible and undesirable, and that conclusion seems to carry across to everyday life.

millennia. Perhaps we should say it's debating whether or not to stay as a monadic. Millennial symbolism is an important driving force here, and such symbolism is distorting election campaigns the world over.

Chapter 8

Constructivist Probability

It is a standard claim of modern epistemology that reasonable epistemic states should be representable by probability functions. I'll call theories which make this claim *classical*, which seems an accurate enough label. That claim is relaxed in this dissertation only to the extent of allowing vague credences, i.e. allowing epistemic states to be vague over a set of probability functions. It is the set of functions, rather than a single function, which represents the epistemic state. However, there have been a number of authors who have opposed this claim. For example, it has been claimed that epistemic states should be representable by Zadeh's fuzzy sets, Dempster and Shafer's evidence functions, Shackle's potential surprise functions, Cohen's inductive probabilities or Schmeidler's non-additive probabilities.¹ Indeed the move to allowing vagueness has grown to some extent from this opposition to orthodoxy.

In this chapter I will argue that many of their motivations can be better captured by what I'll call *constructivist* theories of probability. These theories allow axiomatisations which are virtually identical in their formal structure to classical axiomatisations of probability. In the classical axiomatisation, however, there is reference to an entailment relation. The principle difference is that constructivist theories interpret this as a reference to intuitionist entailment, and classical theories as a reference to classical entailment.

8.1 Motivations for a Constructivist Approach to Probability

There are four main reasons for grounding the axioms of probability theory in an intuitionist entailment relation rather than a classical one. These are: a commitment to verificationism, a commitment to anti-realism, preservation of the axiom of addition, and avoidance of direct

¹ For more details, see Zadeh (1978), Dempster (1967), Shafer (1976), Shackle (1972), Cohen (1977), Schmeidler (1989).

arguments for the orthodox approach. Now some of these will be viewed by some people as bad reasons for adopting the given position, and I have some sympathy with that view. In particular, the verificationist and anti-realist elements of the theory might well be viewed as negatives. These arguments are principally directed at showing that by their own lights, various heterodox theorists would be well advised to adopt the constructivist theory outlined here. For this reason, I think that this theory is the best competitor to the theory developed in the earlier chapters.

A standard objection to classical approaches is that they have no way of representing complete uncertainty. Because of the failures of Laplace's principle of indifference, it can't be said that uncertainty about p is best represented by assigning credence $1/2$ to p . Heterodox approaches usually allow the assignment of credence 0 to both p and $\neg p$ when an agent has no evidence at all as to whether or not p is true. Because these approaches generally require an agent to assign credence 1 to classical tautologies, including $p \vee \neg p$, these theories must give up what I'll call the *axiom of addition*.

Addition: For disjoint A, B : $Bel(A \vee B) = Bel(A) + Bel(B)$.

Where no ambiguity results I'll also use the term 'axiom of addition' to refer to the equivalent rule for probabilities. Now in some writings (particularly Shafer) the grounding for this is openly verificationist. Shafer says that when an agent has no evidence for p , they should assign degree of belief 0 to p . Degrees of belief, under this approach, must be proportional to evidence². In recent philosophical literature, this kind of verificationism is often accompanied by an insistence that only intuitionistically valid deductions are sound arguments.

² This assumption was shared by many of the participants in the symposium on probability in legal reasoning, reported in the Boston University Law Review 66 (1986).

A similar kind of argument is made by Harman (1983). He notes that when we don't distinguish between the truth conditions for a sentence and its assertibility conditions, the resultant logic is intuitionist. And when we're considering gambles, something like this is correct. When betting on p we don't, in general, care if p is true as opposed to whether it will be discovered that p is true. A p -bet becomes a winning bet not when p occurs, but when p becomes assertible. So perhaps not just verificationists like Shafer, but all those who analyse degrees of belief as propensity to bet should adopt constructivist approaches to probability.

To see the point Harman is making, consider this example. We are invited to quote for p -bets and $\neg p$ -bets, where p is *O. J. Simpson murdered his wife*. If we are to take the Californian legal system literally, the probability of that given the evidence is strictly between one-half and one. To avoid one objection, these bets don't just pay \$1 if the bettor guesses correctly. Rather they pay \$1 invested at market rates of interest at the time the bet is placed. The idea is that if we pay x cents for the bet now, when it is discovered that we have bet correctly we will receive a sum of money that is worth exactly as much as \$1 now. Still, I claim, it might be worthwhile to quote less than 50 cents for each of the bets. Even if we will receive \$1 worth of reward if we wager correctly, there is every possibility that we'll never find out. So it might be that placing a bet would be a losing play either way. To allow for this, the sum of our quotes for the p -bet and the $\neg p$ -bet may be less than \$1. As Harman points out, to reply by wielding a Dutch Book argument purporting to show that this betting practice is incoherent would be blatantly question-begging. That argument simply assumes that $p \vee \neg p$ is a tautology, which is presumably part of what's at issue.

Harman's point is not to argue for a constructivist approach to probability. Rather, he is arguing against using probabilistic semantics for ordinary propositional logic. Such an approach he claims would be bound to lead to having an intuitionist logic for the reasons given above. He thinks this would be an error, hence the move to probabilistic semantics is simply an error. Whatever we think of this conclusion, we can press into service his arguments for constructivist probability.

The second argument for this approach turns on the anti-realism of some heterodox theorists. So George Shackle, for example, argues that if we are anti-realist about the future, we will assign positive probability to no future-directed proposition. The following summary is from a sympathetic interpreter of Shackle's writing.

[T]here is every reason to refuse additivity: [it] implies that the certainty that would be assigned to the set of possibilities should be 'distributed' between different events. Now this set of events is undetermined as the future – that exists only in imagination – is. (Ponsonnet, 1996: 171)

Shackle's anti-realism is motivated by what most theorists would regard as a philosophical howler; he regards realism about the future as incompatible with human freedom, and holds that humans are free. The second premise here seems harmless enough, but the first is rather difficult to motivate. Nevertheless, there are some better arguments than this for anti-realism about the future. If we adopt these, it isn't clear why we should 'assign certainty' to the set of possibilities.

Shackle is here assuming that for any proposition p , even a proposition about the future, $p \vee \neg p$ is now true, although neither disjunct is true. Given his interests it seems better to follow Dummett here and say that if we are anti-realists about a subject then for propositions p about that subject, $p \vee \neg p$ fails to be true. Hence we have no need to 'assign certainty to the set of possibilities'. Or perhaps more accurately, assigning certainty to the set of possibilities does not mean assigning probability 1 to $p \vee \neg p$.

The third motivation for adopting a constructivist approach to probability is that it allows us to retain the Kolmogorov axioms for probability, and, in particular, to retain the axiom of addition. This axiom has, to my mind at least, some intuitive motivation. And the counter-examples levelled against it by heterodox theorists seem rather weak from the constructivist perspective. For they all are cases where we might feel it appropriate to assign a low probability to a proposition and

its negation³. Hence if we are committed to saying $Pr(p \vee \neg p) = 1$ for all p , we must give up the axiom of addition. But the constructivist simply denies that in these cases $Pr(p \vee \neg p) = 1$, so there is no counter-example to addition.

The final argument for taking a constructivist approach is that it provides a justification for rejecting the arguments of chapter 3. There I provided a new justification for requiring coherent degrees of belief to be representable by the classical probability calculus. The justification, however, simply assumed classical, rather than say intuitionist, logical reasoning was appropriate. The constructivist has a principled reason for rejecting those arguments. The person who adopts a classical propositional logic, but a non-classical probability logic, has not.

8.2 The Morgan - Leblanc - Mares Calculus

In a series of papers (Morgan and Leblanc (1983a, 1983b), Morgan and Mares (1995)) an approach to probability grounded in intuitionist logic has been developed. The motivation is as follows. A machine contains an unknown set of propositions S , which need not be consistent. $Pr(A, B)$ is the maximal price we'd pay for a bet that S and B intuitionistically entail A . By standard Dutch Book arguments, we obtain axioms for a probability calculus which has some claim to being constructivist. The point of this section is to point out the shortcomings of this approach as a theory of uncertain reasoning from evidence. That is, I point out the implausibility of interpreting the axioms they derive as normative constraints on degree of belief.

The axiomatisations given in the 1983 papers differs a little from that given in the 1995 paper, but the criticisms levelled here apply to their common elements. In particular, the following four axioms are in both sets.

³ Again, the discussion in Shafer (1976, chapter 2) is the most obvious example of this, but similar examples abound in the literature.

$$(C1) \quad 0 \leq Pr(A, B) \leq 1$$

$$(C2) \quad Pr(A, A \& B) = 1$$

$$(C3) \quad Pr(A, B \& C) \cdot Pr(B, C) = Pr(B, A \& C) \cdot Pr(A, C)$$

$$(C4) \quad Pr(A \supset B, C) = Pr(B, A \& C)$$

These four are enough to get both the unwanted consequences. In particular, from these we get the ‘no negative evidence’ rule: $Pr(A, B \& C) \geq Pr(A, B)$. The proof is in Morgan and Mares (1995: 458). Now given the semantic interpretation they have adopted, this is perhaps not so bad. After all, if we can prove A from B and S , we can certainly prove it from $B \& C$ and S , but the converse does not hold. However, for the purposes we have adopted it seems a little implausible. In particular, if C is $\neg A$, it seems we should have $Pr(A, B \& \neg A) = 0$ unless B entails A , in which case $Pr(A, B \& \neg A)$ is undefined.

It shouldn’t be too surprising that we get odd results given (C4). Lewis (1976) shows that adopting it for a defined connective ‘ \rightarrow ’ within the classical probability calculus leads to triviality. And neither the arguments he uses there nor the arguments for some stronger conclusions in Lewis (1986) rely heavily on classical principles. The 1983 papers by Morgan and Leblanc don’t discuss this threat, but the 1995 paper takes it seriously. While Morgan and Mares claim to have escaped the threat of triviality, they seem to have done so only by lowering the threshold.

In intuitionist logic we often take the falsum \perp as a primitive connective, with $\perp \supset A$ a theorem for any proposition A . Hence a set of propositions S is consistent iff it doesn’t entail \perp . Now it seems plausible, at least from the perspective we’ve adopted, to take the following as an axiom.

$$(C\perp) \quad \text{For consistent } B, Pr(\perp, B) = 0.$$

Given consistent evidence, we have no evidence at all that the falsum is true. Hence we should set the probability of the falsum to 0. Given Morgan and Leblanc’s original semantic

interpretation there is less motivation for adopting $(C\perp)$, since S might be inconsistent. The restriction to consistent B in $(C\perp)$ is because I take $Pr(A, B)$ to be undefined for inconsistent B . Morgan, Leblanc and Mares take it to be set at 1. The choice here is a little arbitrary, the only decisive factor seems to be what makes for easier statement of theorems. Intuitionistically, \neg is often introduced as a defined connective, as follows.

$$\neg A =_{df} A \supset \perp$$

Assuming A & B is consistent, it follows from $(C4)$ and $(C\perp)$ that $Pr(\neg A, B) = 0$. Again, from my perspective this is an implausible result. The main purpose of this section has been to show that the Morgan - Leblanc - Mares probability calculus cannot do the work I am wanting a probability calculus to do. That is, it is implausible to regard their $Pr(A, B)$ as the reasonable degree of belief in A given B . Hence the logic they have developed cannot be the constructivist one that I argued in section 1 heterodox theorists should endorse.

8.3 Developing a Constructivist Probability Calculus

The principle motivation for constructivist approaches to probability was a form of verificationism. This should be reflected in a constructivist interpretation of probability sentences, and indeed of sentences about degrees of belief. On a classical approach I interpreted $Bel(A) = 1/2$ as meaning the agent has the same degree of belief in A as they have in a fair coin landing heads if tossed. On the constructivist approach I interpret $Bel(A) = 1/2$ as meaning that the agent has as much evidence for A as they have for the proposition ‘The coin will land heads’. There is a difference between the epistemic attitude they take towards A and the attitude they take to the coin toss. More evidence could come in for A in the sense that they could (at least in some cases) become more confident in A without becoming less confident in its negation. In the case of the coin this is not possible. Any evidence for A , like seeing the coin, or hearing someone say, “The coin landed heads” will be evidence against $\neg A$.

This difference with the classical approach is reflected in how I model probabilistic beliefs. In the classical approach I required that if $Bel(A) = 1/2$ then there were dummy propositions p_1, p_2 such that $A \equiv p_1$ is an element of the model, as is $\neg A \equiv p_2$. I then imposed further conditions reflecting the fact that p_1 and p_2 are modelling exclusive, exhaustive and equally probable propositions. In fact, given these conditions, we can deduce $\neg A \equiv p_2$ from $A \equiv p_1$, and hence it follows that if $Bel(A) = 1/2$ then also $Bel(\neg A) = 1/2$. Since this is not a result constructively that is wanted, I need to change something. The considerations of the previous paragraph suggest that I shouldn't require $A \equiv p_1$ to hold in the model. Rather I should just require $p_1 \supset A$. This is interpreted as meaning that for any evidence we have for p_1 there is matching evidence for A , but there might be more evidence for A to come.

To ease the exposition, I'll simply define a constructivist probability function at this stage as any function from sentences to reals satisfying the following three axioms, with the justification of this description coming later.

$$(CP1) \quad 0 = Pr(\perp) \leq Pr(A) \leq Pr(A \supset A) = 1$$

$$(CP2) \quad \text{If } A \vdash B \text{ then } Pr(A) \leq Pr(B)$$

$$(CP3) \quad Pr(A) + Pr(B) = Pr(A \vee B) + Pr(A \& B)$$

In Appendix 8A I show that any epistemic state which is coherent under the above definition of degrees of belief must be representable by a constructivist probability function. Under the simplifying assumption that all degrees of belief are rational numbers, the proofs are entirely constructive; however, in the general case where we just assume degrees of belief are real numbers I can't get all the results constructively.

The entailment here in (CP2) is read intuitionistically. As noted already, these axioms take probability to be a function from sentences to numbers, whereas, in the classical case, the domain of the function was a set of propositions. That equivalent sentences have the same probability is

a theorem we quickly derive from (CP2). Given that, the axioms as stated are just about independent. (Once we have $0 = Pr(\perp)$ and $Pr(A \supset A) = 1$ the rest of (CP1) follows from (CP2)). In particular, (CP3) is not entailed by the axiom of addition along with (CP1) and (CP2), as it is classically, although the axiom of addition does follow from (CP3) and (CP1).

8.4 Kripke Trees

Classically, probability theory is just a special case of measure theory. A probability function is a normalised measure over a possibility space. Indeed, a function is a probability function iff it can be expressed as a normalised measure over a possibility space. It would be convenient for technical purposes if we could find a similar way of characterising constructivist probability functions. An attempt to do this will be made here, but it isn't yet a totally successful attempt. A measure-theoretic account of constructivist probability functions will be developed which is sound with respect to the axioms given above, but I have no proof that it is complete. That is, any function which can be expressed as a measure of the type I am discussing is a constructivist probability function, but I don't have a proof that all constructivist probability functions can be so expressed.

There is still I hope some interest in this measure-theoretic approach. First, it might be subsequently proven that the semantics provided here are complete with respect to the axioms set out above. Secondly, if it turns out not to be complete, we can always add more axioms to make it complete. And if this happens we might have some reason for thinking that the new axioms are justified. That is, the measure-theoretic semantics might give us a clearer guide as to what should be the coherence constraints on reasonable belief.

I ought to clarify the direction of the argument here. As in the classical approach, I take the definition of quantitative credences to be basic. Whatever coherence constraints there are on credences should be justified in terms of these definitions. In the previous section, I noted that the

three axioms could be so justified. The measure-theoretic account given here is a different attempt to capture the class of coherent functions. To the extent that this approach follows the definitions more closely, the coherence constraints derivable from it seem plausible candidates for real normative constraints on rational agents.

In Kripke (1965), a semantics for intuitionist propositional and predicate logic is developed. I'm only interested here in the semantics for propositional logic. The semantics is similar in some ways to the semantics he had earlier developed for classical modal logic. The semantics is based around what are now known as Kripke trees. These are partially ordered sets with a certain type of valuation on them. It is a little misleading to think of the elements of the sets as possible worlds, and the valuations as saying what is true and false at these worlds. This is misleading because the notion of truth at a world employed in this interpretation is non-constructive.

A better (though still perhaps not entirely accurate) interpretation is to say that elements of the set (what we'll call *nodes*) are knowledge states. So the valuation assigns to each node the set of propositions discovered by that time. We don't say that propositions not discovered by that node are not true there, unless we know that we will not discover the proposition to be true at any later node. The formal definitions are as follows⁴:

A Kripke tree for a propositional language (i.e. set of proposition letters) L is a triple (K, \leq, \Vdash) where (K, \leq) is an inhabited⁵ partially ordered set and \Vdash is a subset of $K \times L$ such that:

$$(1) \quad \forall k, k' \in K, \text{ if } k \Vdash p \text{ and } k' \geq k \text{ then } k' \Vdash p$$


We read $k \Vdash p$ as k forces p . We extend \Vdash to compounds by the following definitions.

⁴ The setting out here follows closely Troelstra and Van Dalen (1987: 75-78).

⁵ Constructively, saying that a set has elements, is inhabited, is stronger than merely saying it is non-empty.


- K1 $k \Vdash A \& B$ iff $k \Vdash A$ and $k \Vdash B$
- K2 $k \Vdash A \vee B$ iff $k \Vdash A$ or $k \Vdash B$
- K3 $k \Vdash A \rightarrow B$ iff for all $k' \geq k$ if $k' \Vdash A$ then $k' \Vdash B$
- K4 not $k \Vdash \perp$

We then introduce \neg as a defined connective, $\neg A =_{\text{df}} A \supset \perp$. If we interpret $k' \geq k$ as k' is subsequent to k then (1) implies that whatever is discovered is not forgotten. (1), K3 and K4 imply that we only discover $\neg A$ when we know that we won't ever discover A . K2 implies that we can only discover a disjunct by discovering one or other disjunct. The following tree shows how we can discover $\neg\neg A$ at a node without discovering A at that node.



At node 0 we have not discovered A . However, we do know that whenever we discover $\neg A$ we'll discover \perp . This follows from the fact that we know we'll never discover $\neg A$. So node 0 forces $\neg\neg A$ without forcing A . So at 0 neither $\neg\neg A \supset A$ nor $A \vee \neg A$ are forced.

The following example shows that $\neg A \vee \neg\neg A$ is not forced at all nodes. At node 2 we have not discovered A and there are no subsequent nodes, so we know we won't discover A . Hence $2 \Vdash \neg A$. So we don't have $0 \Vdash \neg A$. And since $1 \Vdash A$, we don't have $0 \Vdash \neg\neg A$. Hence we don't have $0 \Vdash \neg A \vee \neg\neg A$.



Kripke models are important because for every sequent which is not intuitionistically valid there is a Kripke model with a node which forces all the premises but which does not force the conclusion. The original proof of this in Kripke (1965) was non-constructive; however, it has

subsequently been constructively proven. None of this is at all new to readers familiar with intuitionist logic. However, since at least part of the purpose of this chapter is to promote constructivist approaches to theorists presumably unfamiliar with them, a small background is probably in order.

With the Kripke trees, we can develop a measure-theoretic semantics for constructivist probability. In short, a constructivist probability function is a normalised measure on a Kripke tree. This needs some explaining and some justifying; the explaining first.

Strictly speaking, since Kripke trees can be uncountably large, we need to be more precise than just saying a measure is placed on the tree. Rather, for any tree (K, \leq, \Vdash) and language L we must first define a field F of subsets of K as follows. For any sentence A of L , define K_A to be the set of nodes of K which force A . (Languages, by the way, are customarily taken to have only a denumerable number of sentence letters in them. I'll follow this convention here. From this it follows there are only denumerably many finite sentences of L .) Let F be the smallest set of subsets of K which includes K_A for every sentence A , and is closed under complementation, finite⁶ intersection and union.

Now define a measure m on F to be any function onto reals satisfying the following rules. (D and E are arbitrary elements of F .)

$$(m1) \quad m(\emptyset) = 0 \leq m(D) \leq 1 = m(K)$$

⁶ The definition of F could be extended to ensure it is closed under countable intersection and union, and similarly the definition of measure to ensure it is countably additive. This would of course ensure that an axiom stronger than (CP3) was needed. Since constructive approaches were developed out of a desire to remove completed infinities from theory, perhaps this would be a mistake, but I don't intend to go into this question here.

(m2) If $D \cap E = \emptyset$ then $m(D) + m(E) = m(D \cup E)$

I then define a constructivist probability function as $Pr(A) = m(K_A)$ for any sentence A . A non-constructive proof that probability functions so defined will satisfy the CP-axioms can be easily given. The proof requires the assumption that for any node k , $k \in K_A$ or $k \notin K_A$, equivalently that k either forces or doesn't force A . I could perhaps fix this problem by insisting that the trees be finite. (Intuitionist logic is complete with respect to finite Kripke trees.) However, that would generate problems of its own. For example, there ought be at least one probability function with $Pr(q_i) = 1/i$ for all i , yet this would be ruled out if we said all probability functions were measures on finite trees.

If I weakened (CP3) to just say that, when A and B are provably disjoint, the probability of their disjunction is the sum of their probabilities, then I can give a constructive proof that all probability functions defined by a measure will satisfy the axioms. However, there appear to be good reasons (from the definition of degrees of belief) to insist on the stronger axiom. Perhaps the best solution here is to stipulate that \Vdash is decidable by definition. The justification is that the nodes are not real states but rather representation of states of inquiry. That is, the nodes are individuated by what will be known at each node. Hence we should be able to tell, by virtue of the fact that we can individuate the nodes, what will be known there. An alternative solution is to allow that the metatheory is classical, but I am aiming here to do as much as possible constructively.

So assuming \Vdash is decidable, I'll quickly list the proofs that all the probability functions defined in this way satisfy the CP-axioms. Since $K_{\perp} = \emptyset$, and $m(\emptyset) = 0$, $Pr(\perp) = 0$. Similarly $K_{A \supset A} = K$, so $Pr(A \supset A) = m(K) = 1$. For any A , $m(K_A) \in [0, 1]$, hence $Pr(A) \in [0, 1]$. This proves (CP1).

Assume $A \vdash B$. Hence $K_A \subseteq K_B$. So $K_B = K_A \cup (K_B \cap K_A^c)$. Since $m(K_B \cap K_A^c) \geq 0$, and $m(K_B) = m(K_A) + m(K_B \cap K_A^c)$, $m(K_B) \geq m(K_A)$. Hence $Pr(B) \geq Pr(A)$, proving (CP2). The proof here

would be non-constructive if I wasn't assuming \Vdash is decidable, but this could be fixed if I added to the definition of measure that no subset has greater measure than its superset.

Finally, $K_{A \& B} = K_A \cap K_B$, and $K_{A \vee B} = K_A \cup K_B$ by the definition of Kripke trees, so by (m2) it follows that $m(K_{A \& B}) + m(K_{A \vee B}) = m(K_A) + m(K_B)$. From this (CP3) follows trivially. This requires that $K_A = (K_A \cap K_B) \cup (K_A \cap K_B^c)$, which wouldn't be a legitimate assumption constructively unless I had assumed \Vdash is decidable.

Why might we think that such measures are representations of reasonable epistemic states? In contrast to the classical case, it doesn't make a lot of sense to regard the measure of individual nodes as particularly meaningful. Rather, the figure relevant to each node is the measure of its descendants. For simplicity, define $m'(k) = m(\{k' : k' \geq k\})$ ⁷. This is the probability that we will reach this node in our explorations; that we will discover all the propositions which are forced by this node. This second level probability behaves classically. Note that all classical probability functions are constructivist probability functions, but that the converse is not the case.

The intuitive justification for the measure theoretic semantics is that as we acquire evidence for a proposition, we increase the probability that we will discover that proposition to be true. When we have little evidence for a proposition, the claim is that we don't have much likelihood of discovering it to be true. This does require a rather generous interpretation of evidence; a method for discovering whether or not p is true counts as evidence both for and against p because it increases the likelihood of discovering p and discovering $\neg p$. Such an interpretation is quite natural constructively, but doesn't make a lot of sense classically. So a motivation for the measure

⁷ This might not be defined in some cases. Whenever a set of nodes D is not an element of F we can approximate its measure above as the lower bound on $\{m(E) : E \supset D\}$ and below as the upper bound on $\{m(E) : E \supset D\}$.

theoretic approach can be found within the definition of numerical credences, which as I said above remains the core test.

8.5 Intuitionist Probability

A constructivist probability function can satisfy $Pr(A \vee B) = 1$ without satisfying $Pr(A) = 1$ or $Pr(B) = 1$. To the extent that the assertibility of a proposition is given by its having probability 1, this means that an agent can assert disjunctions without being able to assert either disjunct. Since the denial of this possibility is sometimes taken to be a distinctively intuitionist claim, it seems constructivist probability functions are not intuitionistically acceptable. It seems that it would be more acceptable from this approach to use a probability calculus based, say, on fuzzy logic, such that the probability of a disjunction is the higher of the probabilities of the two disjuncts.

We don't need to be so radical. The intuitionist probability functions can be defined as those constructivist probability functions that satisfy the following axiom, (again A and B range over all propositions):

(CP4) If $Pr(A \vee B) = 1$ then $Pr(A) = 1$ or $Pr(B) = 1$.

I mentioned above that all classical probability functions were constructivist probability functions. However, they will not, in general, be intuitionist probability functions. Indeed, by letting B be $\neg A$, the only classical probability functions which will satisfy (CP4) are those that set $Pr(A) = 0$ or 1 for all propositions; what Lewis calls 'opinionated functions'.

It's non-trivial to specify necessary and sufficient conditions on measures such that they satisfy (CP4)⁸. One clearly sufficient condition (which has some independent motivation) is the

⁸ Part of the difficulty is the role of \Vdash . Usually if we were trying to find necessary conditions for a condition like this we would try to find conditions on triples (K, \leq, m) such that for any \Vdash , (CP4) is satisfied. So, as in the case of the sufficient condition given in the text, the answer will be independent of \Vdash . And this

following. Say a Kripke tree (K, \leq, \Vdash) is grounded iff there is a $k_g \in K$ such that for all $k' \in K$, $k' \geq k_g$. Then a sufficient condition for (CP4) is that $m(k_g) > 0$. If this is the case then the only propositions which receive probability 1 will be those which are forced by every node, and this can't be true of a disjunction unless it is true of one or other disjunct.

I don't want to say that rational agents ought to have their degrees of belief be intuitionist probability functions; (CP4) is, for reasons I'll set out soon, an unreasonable constraint. The only reason I bring it up is to consider this condition on Kripke trees. I claim that it is reasonable to insist that the trees are grounded, but it is unreasonable to require the ground to have positive measure. This restriction is important, as we'll see below, for the interpretation of conditionalisation in this theory.

The nodes represent possible states of knowledge. The paths through the tree represent possible chains of discoveries. But all these possible real-world paths have a common starting point, what we now know. Hence all the paths in the tree should have a common starting point, and hence be grounded.

This doesn't impose any new axioms. To see this, note that we can easily turn an ungrounded tree (K, \leq, \Vdash) into a grounded one (K', \leq', \Vdash') in the following way:

$$K' = K \cup \{k_g\}$$

$$k \leq' l \text{ iff } k \leq l \text{ or } k = k_g$$

$$k \Vdash' A \text{ iff } k \Vdash A \text{ or } k = k_g \text{ and } \forall l \in K: l \Vdash A$$

simplifies the theoretical task; a proposed condition is shown to be necessary by finding a \Vdash such that a breach of that condition combined with that valuation will lead to a breach of (CP4). But when \Vdash is fixed this can't be done. And since m is defined in part in terms of \Vdash , it doesn't make sense to let the valuation vary as the measure stays constant.

By putting the same measure on (K', \leq', \Vdash') as we put on (K, \leq, \Vdash) we can recover the same probability function from a grounded tree that we originally had. So any probability function which can be expressed as a measure on an ungrounded tree can be expressed as a measure on a grounded tree.

The above construction relied on the acceptability of giving the ground measure 0. This can lead to breaches of (CP4). However, these seem perfectly acceptable. To simplify, we'll consider whether it can ever be reasonable to give probability 1 to $A \vee \neg A$ without giving probability 1 to either A or to $\neg A$. Knowing $A \vee \neg A$ is constructively equivalent to knowing A is decidable in principle. And this is weaker than knowing which of these is true. For example, it is constructively acceptable to say that a certain large number, say 5691364391 is either prime or not prime without being able to say which, simply because we know that there is a way of finding out which is true.

So we only can't assert $A \vee \neg A$ when we don't have evidence to say that A will be decidable in principle. It is, however, a little difficult in practice to know how to apply this. The constructivist rules for assertibility were developed with specific applicability to mathematics, and there we have a much tighter definition of 'in principle' than in the real world. To take a simple example, should we say it's decidable 'in principle' whether a Democrat will win the 2200 U.S. Presidency. No one reading this will, I presume, ever know whether this is true. So ought it count as something undecidable? More importantly, can we come up with a rule which determines what is and isn't decidable in principle?

The main reason this is important is that the reductive definition of quantitative degrees of belief only makes sense if it is assumed that lotteries, or at least dummy lotteries, are decidable. I assumed implicitly that one could assign probability 1 to the proposition *Some ticket will win* without assigning probability 1 (indeed while assigning probability $1/n$) to any proposition of the

form *This ticket will win*. So the definition of decidability has to be set so that these lotteries are decidable, but arbitrary future events need not be.

The best way out of this, I think, is to amend the definition of the dummy lotteries. These were already a little fanciful because we assumed the agent knew them to be fair. Let's add the extra assumption that the agent knows them to be fair and decidable. That is, the agent knows that they will, shortly, know the result of the lottery. I can then take a rather strict stance on what is decidable in the real world. In particular any proposition about the unobserved, be it past, present or future unobserved, is not known to be decidable. I am using 'observed' in an odd way here because we can on this usage observe any decidable mathematical proposition. We can observe that is, whatever Brouwer's ideal mathematician could observe.

Despite this restrictive use of decidability in relation to real world events, I leave open the possibility that we can coherently stipulate a proposition to be known to be decidable without it being known whether the proposition is true or false. This is important for getting a reductive definition of quantitative degrees of belief, without which we cannot develop a probability calculus. So even if there aren't any falsifying instances of (CP4) in the real world (that is, if a rational agent in the real world never believes a disjunction to degree 1 without believing one or other disjunct to degree 1) it oughtn't to be a general normative requirement.

8.6 Updating

There are several problems with accurately capturing updating within the constructivist probability calculus. Some of these can be attributed to the fact that updating, as it is usually discussed within the classical literature, essentially uses non-constructive language. Thus, in part, the difficulties we discover can be attributed to the non-constructive character of updating rules. However, it isn't clear that they all can be so avoided, and the impossibility of finding a fully

justified updating rule within the constructivist approach to probability should count against it to some degree.

To see the problem, just consider a simple example. Say an agent has degree of belief $1/3$ in A and $1/3$ in $\neg A$. It is as if we have an urn with n red marbles, n black marbles and n white marbles, from which we know one will be drawn at random. The agent has as much evidence that A as she has that a black marble will be drawn, and as much evidence that $\neg A$ as that a white marble will be drawn. So we develop a fiction in which the drawing determines what really happens; if a white marble is drawn, $\neg A$ happens; if a black marble is drawn, A happens; and we don't have the evidence to say what happens if a red marble is drawn.

Classically, discovering A was taken to have the same effect as discovering a black marble will be drawn. However, it isn't clear why this should hold constructively. After all, the agent doesn't know (even in the fictional model) that the only way for A to be true is that a black marble be drawn. The drawing of a black marble is a sufficient but unnecessary condition for A . So should we treat a discovery that A as being equivalent to finding out that a white marble will not be drawn?

No, and for two reasons. This would seem to have the consequence of saying $Pr(B | A) = Pr(A \& B) / 1 - Pr(\neg A)$, which would imply that in general $Pr(A | A) \neq 1$. If we just add to the above evidence the knowledge that a white marble won't be drawn, we don't have any reason to be certain that A will happen. A red marble might mean that A happens, but it might not. Since $Pr(A | A) = 1$ looks like a pretty good candidate axiom, more care is needed with the semantics.

Secondly, problems arise if we know something about the red marbles. Say we know that half the black marbles and half the red marbles are B marbles (i.e. if they are drawn, B will happen). And we know the other half of the red and black marbles are $\neg B$ marbles. By the above formulae

$Pr(B | A) = 1/4$, which seems absurdly low, the same way that $Pr(A | A) = 1/2$ seemed low. However, there isn't a particularly easy way to fix this problem.

The difficulty is that it seems $Pr(B | A)$ should be set as the greatest lower bound of the possible ratio of B marbles left to total marbles left, after taking this new evidence into account. We know that all the white marbles have been removed, but we don't know what has happened to the red marbles. We can fix the problem for determining $Pr(A | A)$ by saying we know that all the red marbles which really were $\neg A$ marbles have been removed, hence the ratio of A marbles to total marbles must be 1. However, for $Pr(B | A)$ the situation is more complicated. There is the possibility that all the red marbles which are B marbles are also $\neg A$ marbles. So in the worst case scenario, there are n black marbles left, half of which imply B and half of which imply $\neg B$, and $n / 2$ red marbles left, all of which imply $\neg B$. So only $1/3$ of the remaining marbles are B marbles. Hence $Pr(B | A) = 1/3$. Similar reasoning shows that $Pr(\neg B | A) = 1/3$. However, whatever happens to the red marbles, we know they all imply $B \vee \neg B$. So $Pr(B \vee \neg B | A) = 1$. Hence $Pr(\cdot | A)$ is not a constructivist probability function.

The only way to get out of this problem is to define $Pr(B | A)$ using Bayes's rule, however hard it is to give a constructivist justification of this using the same approach we used in chapter 3 for the classical justification of it. We might be able to get a better motivation using the Kripke trees discussed in section 4. The nodes in a Kripke tree were interpreted to be possible states of knowledge. If we discover A we must, therefore, move to one of the nodes at which A has been discovered. So if we amend the tree by eliminating all those nodes at which A is not forced and re-normalise the measure by multiplying through by a constant, we will have a plausible updated tree. And the updated measure of K_B , i.e. the updated probability of B , will be given by Bayes's rule.

There are some problems too with this approach. The first is that it isn't clear why we should re-normalise in this approach. We could re-normalise by something akin to imaging, moving the measure from the deleted nodes to the nearest undeleted nodes. There's no reason why we should do things this way, but on the other hand I can't see a knock-down argument as to why we shouldn't.

The second is that the updated tree will not necessarily be grounded. This is connected to the problem I mentioned above that updating might not be a constructively acceptable concept. Again, it's simpler to see what's going on in the non-probabilistic case. The following sequent is not constructively valid.

$$A \supset (B \vee C) \not\vdash A \supset B \vee A \supset C$$

This follows from the constructive interpretation of \supset . There, $A \supset B$ is interpreted as meaning there is a construction which transforms every proof of A into a proof of B . The above sequent fails because there might be a construction which transforms every proof of A into either a proof of B or a proof of C , but the transformation takes some proofs of A into proofs of B and some into proofs of C .

So we can imagine a rational agent who believes $p \supset (q \vee r)$ without believing either $p \supset q$ or $p \supset r$. Assume that this agent follows the intuitionist rule of never believing a disjunction without believing one or other disjunct. And assume too that the agent follows the plausible updating rule of coming to believe B iff there is some A such that they believe $A \supset B$ and discover A . What should the agent do upon discovering only that p is true? At first it looks like they should come to believe $q \vee r$ without believing either disjunct. But this is misleading, for we cannot say constructively that they only discover p . When they discover p they must discover it by some process. And that process will either be one which is also a discovery of q or is also a discovery

or r . So they will come to believe one of q and r , but without knowing how they have discovered p we can't know which one.

The point of the story is that just saying that an agent discovers a certain proposition to be true without specifying a process of discovery is not constructively acceptable. In the probabilistic case, if the agent updates with respect to any full process of discovery they will, presumably, move to a certain node in the Kripke tree, and hence the updated tree will be grounded. Having our agents update on a proposition without a discovery process is not constructively acceptable.

In sum, the only plausible constructivist updating rule is Bayes's rule. It isn't obvious that this can be justified, implying that there might be no justifiable updating rule. However, this mightn't be a problem if the constructivist can argue that they should not be required to produce an updating rule because such rules use non-constructive language. It doesn't make sense, goes the objection, to ask what a rational agent would do if they had discovered A but had no process of discovering A . On the other hand if they do have a process of discovery they have more information than just A , and this should be used too, and in these cases Bayes's rule seems unproblematic.

8.7 Objections

Despite my promotion of the constructivist approach in the previous section, it is not the approach which I am endorsing in this dissertation. As mentioned above, it is a good candidate for the second best approach to representing rational states of uncertainty. The advantage it enjoys over the precise classical approach is that it has a way of representing complete ignorance. And this, I think, is a large advantage. So I think going constructivist would be a forward step for theorists who object to the classical approach, even in its imprecise form as advocated throughout this dissertation. There are, though, good objections to be raised. I suspect none of these will be good *ad hominen* arguments; they involve, on the whole, rejecting the constructivist program rather than specifically showing there to be flaws in the constructivist approach to probability.

The first type of objection rests on simply rejecting the philosophical motivation. The motivations I mentioned above were verificationism and anti-realism. But neither of these seem at all attractive as philosophical theories. The defeat of verificationism is one of the great successes of twentieth-century philosophy. And while the realism / anti-realism debate has some life left in it, the kind of anti-realism needed to motivate this approach to probability seems much too far fetched.

I used Shackle's anti-realism about the future to motivate a constructivist approach. And if we're anti-realist about the future we can reasonably have degrees of belief in statements about the future, say A , such that $Bel(A) + Bel(\neg A) < 1$. However, if this is plausible for statements about the future it seems just as plausible for statements about the past. There is little justification for requiring agents' degrees of belief about past-directed statements to be a classical probability function. Hence, to motivate a general constructivist approach we have to give up not only realism about the future, but realism about the past. And that just seems implausible.

The second type of objection turns on the plausibility of classical standards of validity. That is, despite all that has been said above, it does seem we can assert $A \vee \neg A$ for any proposition A on any evidence whatsoever. Equivalently, we can assert if $\neg\neg A$ then A for any A again on any evidence. As an objection this is fairly question-begging; it is much like those occasions where someone claims they can show their opponent is wrong by simply stating their opponent's position loudly and clearly and exclaiming, "No one could believe that!" Still, sometimes it is worth doing. And in fundamental questions, where there is disagreement even about what counts as a conclusive argument, it will be difficult to get arguments which at the end of the day aren't question begging.

Ramsey, in *Mathematical Logic* (1926b), thought he could dismiss intuitionism on this ground, and the extent of his argument was little more than a quote from Lewis Carroll.

“It’s very long,” said the Knight, “but it’s very *very* beautiful. Everybody that hears me sing it – either it brings the *tears* into their eyes, or else –”. “Or else what?” said Alice, for the Knight had made a sudden pause. “Or else it doesn’t, you know.” (Carroll, 1871, 306)

Ramsey of course had a conversion to intuitionism (not quite on his death-bed, but in the last year or so of his life), so perhaps arguments founded on Carroll’s jokes are not too secure. For many, it’s just a Moorean fact that sentences $A \vee \neg A$ are all true. Now Moorean facts come in many varieties; some aren’t even facts. There is perhaps a scale with “Moore had two hands” at one end, and “A set has more elements than any of its proper subsets” at the other. Every claim on the scale is *prima facie* plausible. Somewhere down the line the claims stop being acceptable without argument; presumably, this is a little before they stop being true.

There is a slightly less question-begging approach. It is well known that we cannot conservatively add a classical negation operator to constructivist logic. Assume we have a constructivist propositional logic, with \neg as the negation operator. Then $A \vee \neg A$ is not a theorem. However, if we added a classical negation operator \sim , so that we had $\sim\sim A \supset A$, and $(A \supset \perp) \supset \sim A$, we would be able to prove not only $A \vee \sim A$, (as we want) but $A \vee \neg A$, which isn’t wanted. Hence constructivists have to say that the classical interpretation of the connectives is “unintelligible” (Dummett 1977, 11) that they do not understand what the classical connectives mean. But this last claim beggars belief. When we look at the work constructivists have done in classical logic and mathematics it is perfectly clear they do understand the connectives, perhaps better than most of their opponents.

I don’t expect any of the above to change anyone’s position. It is more a statement of why I don’t think a constructivist approach ought, in the end, to be adopted, rather than an argument against it. The only argument for this which isn’t question-begging is the difficulty constructivist

approaches have in dealing with updating. Even here, there is the possibility of an argument that this would be an unfair requirement. So arguably what we have here are differing approaches which are each internally coherent and which are effectively immune to external challenge. This wouldn't be a disastrous result; if it were true, my preference for a classical approach would be just a matter of taste, and the theory developed in this chapter provides a useful alternative for those with different tastes.

Appendix 8A Proof of Soundness of the Axioms

This appendix formally sets out the constructivist definition of degrees of belief, and proves that this definition entails that (CP1), (CP2) and (CP3) are coherence constraints on degrees of belief.

Let $Bel(A)$ be a function from sentences to the degrees of belief on an agent. Let Γ be a finite set of sentences, closed under negation, conjunction and disjunction, such that for all A in Γ , $Bel(A)$ is rational, and y be the lowest common denominator of the values $Bel(A)$ takes. Let P be the set of dummy propositions $\{p_1, \dots, p_y\}$, which are defined such that the agent has no beliefs about any of the p_i . In the classical case I could make this last condition strict; here, I need to employ a primitive notion of disconnectedness.

The agent's beliefs about Γ are coherent iff they can be modelled by K^* , which is a set of sentences closed under (intuitionist) entailment, and satisfies the following conditions:

- (1) For all A, x , $Bel(A) \geq x/y$ iff $\exists S: (S \subseteq P \ \& \ |S| = x \ \& \ (S \supset A \in K^*))$
- (2) For all $S \subset P$, $S \notin K^*$
- (3) $P \in K^*$
- (4) For all i, j , $\neg(p_i \ \& \ p_j) \in K^*$
- (5) For all i, A, B if $(p_i \supset A \vee p_i \supset B) \in K^*$ then $p_i \supset A \in K^*$ or $p_i \supset B \in K^*$

As in the earlier account, for simplicity I sometimes identify a set with the disjunction of its elements. From (3) and (4) it follows that, for all i , $p_i \vee \neg p_i \in K^*$, a fact I use in some of the proofs below. The justification of (5) is the constructive construal of disjunction. The idea is that we can't be able to say in the model that we either have evidence for A or for B without being able to say one or the other. The aim now is to prove that if Bel can be modelled by K^* satisfying (1) to (4), it must be a constructivist probability function, that is it must satisfy (CP1) to (CP3).

I have assumed that $Bel(A)$ is a rational number whose denominator is a factor of y . Hence there are only finitely many values $Bel(A)$ can take; $y + 1$ to be precise. We also know that it takes at least one value (we'll prove soon it takes at most one). So if we can prove that $Bel(A)$ is not equal to y of these possible values, we will have proven that it must equal the other one. This insight allows us to use *reductio* arguments that are not in general constructively acceptable.

Assume $Bel(\perp) > 0$, so $Bel(\perp) \geq 1/y$. Hence there is a p_i such that $p_i \supset \perp \in K^*$. Since this is the same as $\neg p_i \in K^*$, and since $P \in K^*$, it follows by disjunctive syllogism that $P / \{p_i\} \in K^*$, contradicting (2). As $Bel(\perp)$ must take some value, this implies it must be zero.

As $A \supset A$ is a theorem, so is $P \supset (A \supset A)$. As K^* is closed under entailment, $P \supset (A \supset A) \in K^*$. So by (1) $Bel(A \supset A) = 1$.

Assume $A \vdash B$. So $A \supset B$ is a theorem, and hence is in K^* . Assume $Bel(A) = x/y$. So there is an S of size x such that $S \supset A \in K^*$. By *modus ponens*, this implies $S \supset B \in K^*$. So $Bel(B) \geq x/y = Bel(A)$. This proves (CP2), and as $\perp \vdash A \vdash A \supset A$, this completes the proof of (CP1).

The proof of (CP3) meets one early difficulty. The following classically valid inference is not intuitionistically valid.

$$p_i \supset (A \vee B) \vdash p_i \supset A \vee p_i \supset B$$

So I can't say straight away that evidence for $A \vee B$ is evidence for A or evidence for B . However, the following is valid.

$$p_i \supset (A \vee B), p_i \vee \neg p_i \vdash p_i \supset A \vee p_i \supset B$$

Since I already have $p_i \vee \neg p_i$, the inference goes through. So if $p_i \supset (A \vee B) \in K^*$, then $p_i \supset A \vee p_i \supset B \in K^*$, and hence by (5) $p_i \supset A \in K^*$ or $p_i \supset B \in K^*$. Assume $Bel(A \vee B) = x / y$, and $S \supset (A \vee B) \in K^*$, with $|S| = x$. Then for all p_i in S , $p_i \supset A \in K^*$ or $p_i \supset B \in K^*$. For all p_i not in S , $p_i \supset A \vee B$ cannot be in K^* or else (1) would be breached.

Let S_A be the set of p_i such that $p_i \in S$ and $p_i \supset A \in K^*$, with S_B defined similarly. Since S_A is the largest subset S' of P such that $S' \supset A \in K^*$, so $Bel(A) = |S_A| / y$. Similarly $Bel(B) = |S_B| / y$. Further $(S_A \cap S_B) \supset (A \& B) \in K^*$, and this will not be the case for any larger set, again because if it were (1) would be breached. So $Bel(A \& B) = |S_A \cap S_B| / y$. Finally, because of the results of the last paragraph, $S = S_A \cup S_B$, so $Bel(A \vee B) = |S_A \cup S_B| / y$. In general, for finite decidable sets, $|S_A| + |S_B| = |S_A \cup S_B| + |S_A \cap S_B|$, and this proves (CP3).

Chapter 9

Vague Decision Theory

9.1 Introduction

Part 1 had two central aims: the interpretation of probability sentences and the discovery of rules for reasoning under uncertainty. In the literature, the latter are usually derived by looking at rules for decision making under uncertainty. There are good reasons for thinking this gets matters the wrong way round. So here, in part 2, I move onto decision making under uncertainty as an application of the rules of reasoning already derived. This chapter looks in general at rules; the following two chapters apply my conclusions to some economic problems discussed by Keynes.

Although I call what I'm doing 'decision theory', strictly it is only a portion of a proper decision theory. The central problem with Dutch Book arguments is that they confuse use-value with exchange-value. The same problem I fear pollutes most decision theory. What I'll really be investigating here is the theory of qualitative use-valuation. But that's an ugly name, and since that theory is often called 'decision theory', I'll stick with the ordinary labelling. When I get to applications, the distinction between the two fields will need to be more carefully observed.

The reference to 'qualitative' is important here. I take as given the quantitative use-value of various gambles, and try to derive their relative value. This looks like it should be a triviality: ϕ is more valuable than ψ iff its quantitative value is higher¹. However, the values I take as given are interval-valued, or perhaps set-valued, depending on how we develop the theory. So there's no easy bridge available between quantitative and qualitative theory. Indeed it's unclear whether anyone has developed a bridge that meets some rather minimal coherence requirements. For that

¹ I'll use Greek letters throughout to refer to gambles. I won't define precisely what I count as a gamble, because this is fairly standard in the literature. However, following Ramsey, I regard not only betting slips, but stocks, cars and dollar bills as gambles.

reason, the purpose of this chapter will be largely negative. I will examine many proposals for getting from quantitative (but non-numerical) values to quantitative orderings or practical proposals, and show where they fail. I'll then propose a relatively modest theory that seems at least plausible.

As is usual in the literature, I'll talk about an agent having a choice between various gambles. In order to avoid the problems which brought down the Dutch Book argument, I'll make the following assumptions. Whenever an agent has a choice between a set of gambles, unless otherwise stated, these conditions hold:

- (i) All gambles are to be settled immediately after the choice is made. That is, if an agent has a choice between a p -bet and a q -bet, then immediately after the choice is made, the agent will discover whether their chosen bet wins, and receive the appropriate payout. As a corollary to this, there is no possibility of retrade, and no alternative trades available.
- (ii) The marginal utility of money is constant and independent of the outcome of various bets.
- (iii) The agent places no value on the discovery of whether a particular bet wins or not.

Some of these conditions are inapplicable in some of the circumstances in which we want to apply decision theory. For example, the failure of (i) and (ii) is important when making investment decisions in an inflationary environment. However, it simplifies exposition enormously to assume these conditions hold unless otherwise stated. Condition (i) holds even for decision making in a decision-tree. What I require there is that the choices in the tree are exhaustive, and that bets are settled whenever an 'end-node' is reached.

With these conditions assumed, I can state the central question. Given an agent's (imprecise) degree of belief in each of the relevant propositions, when should an agent trade ϕ for ψ ? This divides into two questions: when is trade permissible, and when is it rationally required? I will also be interested in some associated questions, such as determining which choices are

permissible (or mandatory) from a set of available gambles. The ‘central question’ contains a deliberate asymmetry between ϕ and ψ . I don’t want to rule out theories which say it’s permissible (mandatory) to not trade ϕ for ψ , but not permissible (mandatory) to trade ψ for ϕ . These theories privilege what the agent currently holds. I don’t think such theories work, but this needs argument rather than assumption, and I don’t think this privileging of what an agent holds is an argument against these theories.

The agent is reasonable, so their epistemic state can be represented by a set \mathcal{P} of probability functions. Each of these functions will assign to each gamble a numerical value. Precisely how this is done will depend on how one resolves Newcomblike problems. I won’t buy into this argument, but I’ll assume that it has a resolution². And, since my assumptions mean I can restrict my attention to use-value, I’ll assume that the value, according to Pr , of a p -bet is $Pr(p)$.

For a bet ϕ , \mathcal{P} will determine a range of *expected* values for ϕ : $[l_\phi, u_\phi]$. I’ll assume Continuity, as defined in 7.1.2, holds, so the range of expected values will always be an interval. I’ll always notate it as if it is a closed interval, but there is no reason to assume this. My use of closed intervals is just for simplicity³. It’s very important to remember that $[l_\phi, u_\phi]$ is not the range of possible payouts for ϕ ; that range will usually be considerably wider, and need not be an interval. I am not interested in what the agent thinks ϕ might pay, rather in, roughly, what she thinks ϕ can be expected to pay. If her degrees of belief are all precise then, whatever the range of payouts of ϕ , l_ϕ will equal u_ϕ .

² For my purposes the solution to this problem Lewis gives in his (1981) would work, but any of the similar alternative theories he mentions (such as those due to Sobel, Skyrms and Gibbard and Harper) would also suffice.

³ As is the assumption that l_ϕ and u_ϕ are determined precisely. Because of higher-order vagueness they too will be vague, but that need not detain us.

Decision theories which allow for imprecise credences fall into two broad categories: structured and unstructured. Unstructured decision theories say we can determine the relative merits of ϕ and ψ by just looking at l_ϕ , u_ϕ , l_ψ and u_ψ . Structured decision theories say we need to look at more; in particular, we need to compare the values ϕ and ψ according to particular members of P . The first three theories I'll look at are unstructured; it can be concluded from the way they fail that no unstructured decision theory is plausible.

The bulk of this chapter is negative; I show why a glut of solutions to our problem given in the literature fail. Often I will refer to the advice these solutions give to a contestant in the Monte Hall Problem (MHP). This is formally equivalent to the TPP described in chapter 3, but is perhaps more enjoyable to think about. A contestant on a game show faces three doors, call them a , b and c , with a car behind one and worthless prizes behind the other two. She knows the prize has been allocated by a fair chance mechanism. She chooses a door, and then the host shows her that there is no prize behind one of the doors she hasn't chosen. She knows the host will show her a door, and she knows the host will choose a door to show her that doesn't contain the prize. She is ignorant, however, of the host's procedure for choosing which door to show should she have originally chosen correctly. For convenience, I'll use $\alpha\beta$ ($\alpha, \beta \in \{a, b, c\}$) to refer to the proposition that the car is behind door α and the host shows door β . For a small cost (either in dollars or regret) she is given the chance to change her choice to the other door which remains closed. What ought she do? In chapter 3 I concluded that her degree of belief in 'I originally chose correctly' (call this p) should go from precisely $1/3$ to a vague interval, possibly as large as $[0, 1/2]$. For now I'll assume that her degrees of belief do become that vague; I'll discuss the plausibility of this for the case as described at the end of the chapter.

There is something odd about this epistemic state. The contestant knows that whatever the host does, her attitude towards p will go from being precise to being vague. Fortunately it will always become vague over an interval including the original precise degree, but the interval is guaranteed

to grow. As this causes problems for many of the theories which follow, it might be wondered whether such odd states can be ruled out as unreasonable.

The answer is they cannot, at least on pain of ruling out all vague states as unreasonable. Seidenfeld (1994) shows that on some simple assumptions⁴, the requirement that states be immune to what he calls *dilation* is equivalent to the requirement that states be precise. Let \mathcal{P} be a set of probability functions, and let $\min(h | e)$ and $\max(h | e)$ be defined as the minimal and maximal values respectively of $Pr(h | e)$ for $Pr \in \mathcal{P}$. Let $\Pi = \{p_1, \dots, p_n\}$ be a partition of e . That is, the elements of Π are pairwise disjoint, and their disjunction is e . Then \mathcal{P} is dilated by Π with respect to h and e if for all i , $\min(h | p_i \ \& \ e) < \min(h | e)$ and $\max(h | p_i \ \& \ e) > \max(h | e)$. A set \mathcal{P} is subject to dilation if there is a h , e , and Π such that \mathcal{P} is dilated by Π with respect to h and e . Since requiring that \mathcal{P} be immune to dilation amounts to insisting that \mathcal{P} be a singleton (or satisfy some other even more implausible constraints so that Seidenfeld's assumptions fail) I don't think that requirement can be plausibly imposed. Hence we must learn to live with the decision theoretic consequences of dilation, and the MHP brings out the problem nicely.

To head off possible objections, I will be assuming something like the principle of conglomerability I attacked in chapter 3. Since I will be restricting the scope of the principle to finite partitions this isn't necessarily inconsistent, but it might seem unmotivated. I adopt the following rule:

Restricted Conglomerability

Let ϕ and ψ be bets, such that it would be rationally mandatory for an agent to trade ϕ for ψ were she to learn p , or were she to learn $\neg p$. Then it is rationally mandatory for the agent to trade ϕ for ψ .

⁴ Basically that there are some probabilistically independent propositions.

Now some theorists deny even Restricted Conglomerability because of Prisoners Dilemma and Newcomb Problem cases. I don't want to get into this debate, so I'll just note that I won't use the rule in a way that ought offend such theorists.⁵

In any case, this isn't the objection which concerns me most. What does concern me is the possibility that the arguments in chapter 3 against general conglomerability apply equally here. My only 'argument' for Restricted Conglomerability is its intuitive plausibility. Given this it seems to me quite vulnerable to objections to similar principles supported by the same intuition. I really don't have a response here, except to say that decision theory doesn't seem to get off the ground without a rule like this and giving it up would mean giving up many theories we find attractive. No one has put forward an objection like this in the literature, so perhaps it is best to wait until such objections are made rather than shadow boxing. I think the possibility for objections to the theory of this chapter to work is much greater than for any other part of the dissertation, and this assumption might be one of the weak points.

9.2 Unstructured Decision Theories

9.2.1 *Global Dominance*

Hájek (1998) discusses without endorsing a decision rule called global dominance. This says that it is only rationally compelling to trade ϕ for ψ when $l_\psi > u_\phi$. It isn't made clear, but presumably whenever $u_\psi > l_\phi$ it is rationally permissible to trade. There is a simple counterexample to this

⁵ As a rule like conglomerability is usually needed to justify the idea that we ought to value gambles by their expected utility, objections to it on the ground that it conflicts with the verdict of utility considerations in Newcomblike cases seem to me implausible. Recently, Norton (1998) has argued that we shouldn't accept the verdict of conglomerability in the two-envelope paradox because it conflicts with expected utility considerations. Well, he's right that we shouldn't accept all of conglomerability's verdicts here, but that's because they are inconsistent, not because of the clash with expected utility. Again, without some form of conglomerability there is no motivation for adopting a rule like 'maximise expected utility'.

approach. Let ψ be the bet $\phi + \$\varepsilon$, where ε is some small amount of money such that $l_\psi + \varepsilon = l_\psi < u_\phi$. That is, in any circumstance where ϕ pays $\$m$, ψ pays $\$(m + \varepsilon)$. Clearly here it is rationally compelling to trade ϕ for ψ , however the global dominance rule does not require this.

9.2.2 *Maximin*

Gilboa and Schmeidler (1993) advocate a maximin decision rule. I have already given grounds for rejecting their updating rule⁶, but those objections didn't touch their decision rule. The rule is that it is rationally compelling to trade iff $l_\psi > l_\phi$, and rationally permissible to trade iff $l_\psi \geq l_\phi$. While this rule doesn't give any particularly counterintuitive results for static cases, it seems to do badly in dynamic settings. Of course it wasn't designed to be used with Conditionalisation, so the objection I'm running isn't directed at any particular theorist, just at its possible use with the Bayesian updating rule.

Consider again the MHP with the completely ignorant contestant. Initially $Bel(ab \vee ac) = 1/3$. Hence according to the maximin rule the contestant will gladly buy a $(ab \vee ac)$ -bet for 25 cents. Assume this trade is made. After the host shows the contestant a door, any door will do, the expected value of this bet will now be vague over $[0, 50c]$. Hence by the Maximin rule, she will sell the bet for 20 cents, incurring a sure loss. Hence the Maximin rule, when combined with Bayesian updating, leads to dynamic incoherence.

9.2.3 *Maxi*

This problem could be avoided by adopting a decision rule I call *Maxi*. This says that ψ is strictly preferred to ϕ , i.e. trade is rationally compelling, iff $l_\psi > l_\phi$ and $u_\psi > u_\phi$. Trade is rationally permissible iff $l_\psi \geq l_\phi$ or $u_\psi \geq u_\phi$. No one to my knowledge has endorsed *Maxi* in the literature, but

⁶ In chapter 3.

since it is such an obvious weakening of Maximin and other such rules which have been endorsed, it is worth some discussion.

Although there are no simple examples where Maxi gives counterintuitive results, it is in conflict with conglomerability in some hoked-up examples. If one was committed to Maxi, I suppose it could be said that these were arguments against the sure-thing principle rather than Maxi; or alternatively, that in such bizarre examples we can't expect standard rules to apply. I don't think either of these replies works, but I mention them to note that my objections to Maxi are weaker than my objections to other rules.

Say an agent's degrees of belief are determined by the family of probability functions satisfying the following criteria:

- (i) $0.2 \leq Pr(p \mid r) \leq 0.6$
- (ii) $0.1 \leq Pr(q \mid r) \leq 0.5$
- (iii) $0.3 \leq Pr(p \mid \neg r) \leq 0.7$
- (iv) $0.2 \leq Pr(q \mid \neg r) \leq 0.6$
- (v) $Pr(p) = 0.35$
- (vi) $Pr(q) = 0.4$

It can quickly be seen that none of these conditions are redundant by considering functions like Pr_1 , defined as follows. $Pr_1(p \mid r) = 0.2$; $Pr_1(p \mid \neg r) = 0.7$, $Pr_1(r) = 0.7$, $Pr_1(q \mid r) = Pr_1(q \mid \neg r) = 0.4$. Similar functions show the other six bounds given in the inequalities are non-redundant. Given this epistemic state the value of a p -bet will be precisely 35 cents, and the value of a q -bet precisely 40 cents. However, if the agent were to discover r , the value (in dollars) of a p -bet would be vague over the interval $[0.2, 0.6]$, and that of a q -bet vague over $[0.1, 0.5]$; that is a p -bet would be more valuable, according to Maxi, were the agent to discover r . Similarly if the agent were to discover $\neg r$, the value of a p -bet would go to $[0.3, 0.7]$ and of a q -bet would go to $[0.2, 0.6]$. Again by Maxi, the p -bet would be more valuable.

Hence in these circumstances, Maxi gives the result that a q -bet is more valuable than a p -bet (by 5 cents), however if either r or $\neg r$ were found to be true, it would become the case that a p -bet would be 10 cents more valuable than a q -bet. That is, Maxi is in breach of the conglomerative principle I have adopted. Given that the problem with Maxi is that it is too strong, in the sense that it cannot be that all of the trades which are rationally compelling according to Maxi are really compelling we can draw a more important conclusion. There is no rule expressed purely in terms of l_ϕ , u_ϕ , l_ψ and u_ψ which is stronger than Global Dominance but weaker than Maxi. Yet I've shown that any acceptable rule must be stronger than Global Dominance and weaker than Maxi. Hence no acceptable rule can be expressed purely in terms of l_ϕ , u_ϕ , l_ψ and u_ψ .

As a special case, the Horvitz-style decision rules advocated by Strat (1990) and Jaffray (1994) are incoherent. These advocate that for any bet ϕ we evaluate its expected worth $E(\phi)$ according to this rule.

$$E(\phi) = \rho l_\phi + (1 - \rho)u_\phi. (\rho \in [0, 1]).$$

The operator ρ is an optimism / pessimism operator. The more optimistic we are the higher ρ will be. Since we now have a numerical utility for each bet, we can simply choose the bet with the higher utility. Of course this approach is stronger than Maxi, so if Maxi is too strong, so is this approach. Here the fact that the counterexamples to Maxi are so artificial becomes important, because Strat and Jaffray are not, it appears, aiming to discover the ideal decision rule, but rather trying to find a rule which can be implemented efficiently and gives results which are usually correct. Until an example is found in which the recommendations of this approach are implausible despite the example being realistic enough, their approach might be well-suited to the task they have set themselves.

9.3 Levi's Rule

For the subsequent rules I'll be discussing, I need to look more closely at the structure of the expectation values, not just at their upper and lower bounds. For any bet, say ϕ , and any element Pr of \mathcal{P} , there is a numerical expectation value of ϕ , which we'll call $E_{Pr}(\phi)$. In a completely general theory, utilities as well as credences would be allowed to be vague, but I'll stick to the simple case of assuming precise utilities.

All the subsequent rules I discuss have the property that if for all Pr in \mathcal{P} $E_{Pr}(\psi) > E_{Pr}(\phi)$, then ψ is strictly preferred to ϕ . That is, it is rationally compelling to trade ϕ for ψ . How the rules differ is in what can be done when neither bet is strictly preferred to the other in this sense. For convenience, I'll simply define strict preference to hold between two bets ψ and ϕ iff $E_{Pr}(\psi) > E_{Pr}(\phi)$ for all Pr in \mathcal{P} . This reduces the scope of discussion to bets such that neither is strictly preferred to the other. I will say in this case that the bets are *almost indifferent*. On pain of inconsistency it can't be said that almost indifference implies indifference. This is because almost indifference is intransitive whereas indifference, at least as usually defined, is transitive.

Levi's Rule is that when ϕ and ψ are almost indifferent we should choose the bet which has the highest minimum payout (Levi 1974, 1980, 1986). This minimum payout is referred to as the 'security level' of the bet. I'm keeping with Levi's terminology in referring to choices rather than permissible trades; the translation back into terminology I've been using is usually trivial. He doesn't mean by this that we ought choose ψ iff $l_\psi > l_\phi$. Rather he is referring back to the actual payouts of ϕ and ψ and advocating choice of the bet with the highest possible minimum return, or as he puts it security level. As he notes, when applied to three-way choice this implies violation of the rule of independence of irrelevant alternatives. That is, under his rule it can be rationally required that ϕ be chosen in a pair-wise choice from $\{\phi, \psi\}$, but also required that ψ be chosen in a choice from the set $\{\phi, \psi, \chi\}$. Since he regards the analysis he offers as "impeccable" (1974:

411) he concludes that the rule of independence of irrelevant alternatives must be mistaken in some way.

It's not too surprising that this rule would have to go under such an analysis. After all we can regard each of the *Pr* as a voter which voices an opinion about which choice is best, and then the overall choice becomes the well-known social choice problem of aggregating preferences. Arrow's theorem says that no aggregation rule can satisfy the following four constraints, here explained for voters who are probability functions⁷:

- (1) *Pareto*. If ϕ is strictly preferred to ψ in the above sense ϕ will be chosen from $\{\phi, \psi\}$.
- (2) *Collective Rationality*. The rule determines a preferred option no matter what the various *Pr* functions say about ϕ and ψ .
- (3) *Non-Dictatorship*. There is no *Pr* function whose choice is followed no matter what the other functions say.
- (4) *Independence of Irrelevant Alternatives*. The choice between ϕ and ψ should not depend on what other options are available.

Levi's Rule is committed to (1), (2) and (3), hence it would be inconsistent if it satisfied (4). However, there are good grounds for preserving (4). Of course, there are good grounds for keeping each of these rules, so this argument will necessarily be less than completely compelling. I suspect the strongest argument for (4) is its intuitive plausibility; any attempt to explain this plausibility will sell it short. Nevertheless, I'll try.

Assume an agent, say Lenny, does not satisfy (4). For example, he chooses ϕ from $\{\phi, \psi\}$, but chooses ψ from $\{\phi, \psi, \chi\}$. Assume now he has a choice between $\{\phi, \psi, \chi\}$, but the choice dynamics are as follows. First, he has to specify whether he wants χ or not, and if not he has to

⁷ Arrow's Theorem is set out in Arrow (1963). The setting out here closely follows Hausman (1991).

say whether he wants ϕ or ψ . Lenny's preference is, *ex hypothesi*, to choose ψ , but he can't carry out this choice. Presumably he will reject χ at the first stage, then he will face a choice between ϕ and ψ . And here he is forced by his own preferences to choose ϕ . Levi (1987) in response to this argument claims that Lenny could have adopted at the start a strategy to choose ψ . Hence, at the second stage he will just have to follow his strategy rather than to make a decision about whether ϕ or ψ is preferable. But now the original objection can be restated in a different way. Surely it's a problem for a decision-rule if the only way to consistently implement it is to ignore its recommendations at various stages. Alternatively, it might be argued that the amendment to the rule to allow strategic choice in this way constitutes a rescission of the original rule and substitution of a new rule. The basis for this argument is simply that, according to the amended rule, at times agents are required to act in the opposite way to how they were required to act under the old rule.

Levi tries to minimise this difficulty by saying that it is an ineliminable feature of what he calls 'unresolved conflict'. The problem is that he seems to rely here on some equivocations about what would count as a resolution of a conflict. This leads to a problem, I fear, at the core of his lexicographic approach. Levi thinks that we can have a hierarchy of 'values', such that if we can't decide between two options using the most important value, we can use lower values to resolve it. That's essentially what is being applied here, with expected value being the highest value, and security levels the next. When it is allowed that each of these values might issue non-linear verdicts (they might allow us to be unresolved and not just indifferent between choices) the lexicographic approach hits problems. The problem is essentially that he seems to be committed to saying that some decision making contexts involve a conflict which is essentially unresolved, while at the same time saying that there is a resolution of these conflicts!

Here's an example he gives. Jones, an office manager, has to hire a new worker to do typing and stenographical work. There are three applicants: Jane, Dolly and Lilly. The applicants take tests

in typing and stenography. On the typing test their scores are 100, 91 and 90 respectively, on the stenography test the scores are 90, 91 and 100. So Jones has a dilemma; does he hire the best typist, or the best stenographer, or perhaps someone moderately good at each?

Levi suggests that there are in fact a continuum of tests Jones could apply. For each $\beta \in [0, 1]$ we can work out a candidate's β -score as $\beta x + (1 - \beta)y$, where x is their typing test score and y their stenography score. For each β test there corresponds an argument for selecting the applicant with the highest score on that test. These arguments will often conflict, as in fact they do here. Some tests favour Jane, and some favour Lilly. Since, however, none favour Dolly she can be ruled out. Now Jones is a liberal, but to a degree: he favours using affirmative action criteria to choose a candidate when the continuum of β -tests have failed to be decisive. The affirmative action criteria support ranking the applicants as follows: Dolly, Jane and Lilly. Since Jane is the highest ranked of the candidates left (not ruled out by the β -tests), she gets the job.

But there's a twist to the tale. Just as he's about to tell Jane she has the job, he finds Lilly has withdrawn her application. Now he has to choose between just Jane and Dolly. And since on some β -tests Dolly is now the best of the applicants (where $\beta < 0.1$) she isn't ruled out by those tests. Hence Jones has to make a decision between Jane and Dolly on affirmative action grounds, and *ex hypothesi* Dolly wins. So Lilly's withdrawal means that Dolly now gets the job over Jane.

Levi notes that most decision theorists would demur here. After all, Jones, a poster-boy for his decision-theory, has just violated what we're calling independence of irrelevant alternatives. Here's his defence:

When Jones chooses Dolly, this does not reveal that he thinks Dolly is at least as good as Jane for the job. Jones is in conflict as to who is better, all things considered. He chooses Dolly because in the face of such conflict among the values to which he is committed, he invokes

considerations which otherwise would not have counted for him. When he contemplates the three-way choice, hiring Dolly is ruled out because of his values. This does not mean that his values have changed or that he has inconsistent values. *Hiring Dolly is neither better nor worse than hiring Jane in the two-way choice.* The same remains true in the three-way choice. This example illustrates an important difference between resolving a conflict so that one can choose for the best and failing to resolve a conflict. In the latter case, some consideration which otherwise would not be taken into account is used to provide counsel as to what to do when one cannot choose for the best, all things considered. (1986: 34, my italics)

But this is inconsistent with his description of Jones's motivation. Jones has made a value commitment to hiring on the basis of affirmative action when the β -tests are inconclusive. This is why it can be deduced from his general principles (including his 'tie-breaker' principles) that he will hire Dolly. It is worse, given his principles, to hire Jane over Dolly in the two-way choice, contra what is said in the italicised sentence. Levi wants here to have it both ways; Jones's affirmative action commitment is supposed not to be a value of any kind, so that it wouldn't be against his values to hire Jane over Dolly, but it can at the same time be used 'to provide counsel as to what to do'. It is rather hard to see how this is consistent. To paraphrase Ramsey, if Jones can't say what his choice is, he can't say it, and he can't whistle it either. If two options really are incommensurable, there can't be a reason for choosing one over the other; that just would show that they weren't really incommensurable to start with.

As a footnote to all this, at (1986: 82) he says that rational agents may have a hierarchy of value commitments. This seems to suggest that he favours saying Jones's commitment to affirmative

action is a value, in which case the italicised sentence is simply false, so his general defence here fails.

More difficulties can be made for Levi's decision theory. Assume we have the following test scores for the applicants.

	Typing	Stenography
Tom	100	90
Dick	90	100
Harry	89	99

We have the following affirmative action ordering: Harry, Tom, Dick. If we adopt Levi's rule, we will choose Tom for the position. Dick's scores dominate Harry, so Harry can't pass any of the β -tests. However, both Tom and Dick pass some, so the affirmative action test applies, and Tom is chosen. Now assume that instead of choosing one applicant for a position we have to choose two. We could assume that Tom will be chosen, leaving a two-way choice between Dick and Harry for the final position, which presumably goes to Dick.

It might be thought more efficient, however, to decide whom it would be worst to give the position, and hence offer jobs to the other two. The only plausible way to do this is simply to reverse our tests. So at the first stage we'll look at who's worst on all β -tests, as this is our main criterion. If there is more than one person who is worst according to some β -test, we'll look at who does worst by the affirmative action criteria among these. If we apply this method we find that the worst person to give the job to would be Tom! The only people who are worst according to some β -test are Tom and Harry, and Tom is further down the affirmative action list than Harry. So there are two absurd results: the best person to give the job to is also the worst, and we get different results to the question of which two people we should hire depending on whether we

look for the best two candidates or the worst. For the reasons indicated above, I am unimpressed with Levi's assertions that choices on the basis of 'tie-breaker' principles are not real preferences.

In summary, not only does Levi's rule give counterintuitive results, it rests on a methodology which is suspect because of this equivocation. To add to the counterintuitive results, briefly note the problems Levi's rule has with the MHP. It has just the same problem that the Maximin rule has. Initially $Bel(ab \vee ac) = 1/3$. Hence according to Levi's rule, a contestant will gladly buy a $(ab \vee ac)$ -bet for 25 cents. Assume this trade is made. After the host shows the contestant a door – any door will do – the expected value of this bet will now be vague over $[0, 50c]$. Hence if the contestant is offered 20 cents for her bet, by Levi's rule she will look to security levels. And on these, the 20 cents does best, no matter how we formulate the security rule⁸. Hence by Levi's rule, she will sell the bet for 20 cents, incurring a sure loss. Hence Levi's rule, when combined with Bayesian updating, leads to dynamic incoherence. So there are both theoretical and practical grounds for rejecting Levi's rule.

9.4 Conservatism

The rule I am calling Conservatism is perhaps the dominant decision-theoretic rule amongst Bayesians who allow degrees of belief to be vague. For endorsements of it, see for example Williams (1976) or Seidenfeld (1984) and the references contained therein. The rule is that it's rationally permissible to trade ϕ for ψ iff ψ is strictly preferred to ϕ . As noted above, the rule is asymmetric. There are circumstances in which it is impermissible to trade ϕ for ψ , and impermissible to trade ψ for ϕ . This is an oddity but not an inconsistency. If it was the worst that could be said for the rule it wouldn't be much of an objection. There is, however, a stronger objection.

⁸ I.e. in terms of maximising minimal expected returns or maximising minimal actual returns.

Assume a Conservative is holding ϕ , and ψ is a bet which is almost indifferent to ϕ . Further assume that $\phi + \$10$ is strictly preferred to ψ . The following is a simple-minded objection to Conservatism which doesn't work; I include it to distinguish it from an objection which does work. Assume the only trades which are possible are to swap ϕ for ψ , and, if that swap is made, to swap ψ for $\phi + \$10$. It would clearly be in the agent's best interests to make each of these swaps, but since they are a Conservative they can't make the first swap, hence Conservatism is an irrational rule. The decision-tree is set out in Figure 1.




Figure 9.1

Here's what goes wrong with this objection. When considering the first swap, the Conservative won't be comparing ϕ and ψ ; rather they will be comparing holding ϕ with the possibility of having a choice between having ψ and having $\phi + \$10$. If they had the latter choice, they would choose $\phi + \$10$, hence the original choice is between ϕ and $\phi + \$10$. That isn't much of a choice at all, they will clearly choose the $\phi + \$10$. That is, it is consistent with the Conservative rule to accept both trades.

So this objection fails because it relied on a too simplistic Conservative rule. However, a similar objection can succeed. Alter the payout of accepting both trades to $\phi + \$5$, and assume this is strictly preferred to ϕ , but almost indifferent to ψ . Now the initial choice is a choice between holding on to ϕ , and having the choice between holding ψ or trading it for $\phi + \$5$. The Conservative knows if they have that choice they will hold onto ψ . So now the initial choice

reduces to a choice between holding ϕ and trading it for ψ . Again, the Conservative here prefers to hold ϕ . But this is absurd. Whatever we should end up with in this circumstance, it isn't ϕ , as there is some other option strictly preferred to it. It might be noted that the use of decision-trees in this argument, as opposed to the flawed argument given above, is entirely standard. See Raiffa (1968).

There are two ways out of this problem for the Conservative, neither of them particularly attractive. The first is to make the move Levi makes above, to say that an agent should adopt a strategy for getting through a decision-tree and refuse to reconsider it at later stages. The above objections to that move still apply. The other move is to deny the following rule for reducing complex bets to simple bets.

Reduction: If $C(\underline{\beta}, \chi) = \delta$ for any $\delta \in \{\beta, \chi\}$, then $C(\underline{\alpha}, (\underline{\beta}, \chi)) = C(\underline{\alpha}, \delta)$.

To explain the notation, by $C(\underline{\alpha}, \beta) = \alpha$ I mean that in a choice between holding α and trading it for β , it is rationally compelling that α be chosen. The underlining on α indicates that α is what is currently held; this is important because by the Conservative's lights $C(\underline{\alpha}, \beta) = \alpha$ and $C(\alpha, \underline{\beta}) = \beta$ is consistent. $C(\underline{\alpha}, (\underline{\beta}, \chi)) = \delta$ ($\delta \in \{\alpha, (\beta, \chi)\}$) means that the choice between holding α and trading it for β with the knowledge that this can in turn be traded for χ . Note that I don't assume $C(\underline{\alpha}, \beta)$ is always defined.

I don't have any particularly strong arguments for *Reduction*, but it does have a high degree of intuitive plausibility. It is hard to see what other approach could be taken. As was shown in chapter 3, there is a close relationship between adopting *Reduction* as a decision-theoretic principle and adopting *Addition* as a constraint on credences. Since I have argued that *Addition* is a constraint, because it follows from the Equivalence Analysis, I have a justification for *Reduction*. If anyone thinks it is possible to justify avoiding *Reduction* and hence can avoid this problem I might not have much of a reply. I don't know of any such justification, and I can't see how it could be

intuitively plausible, but I'm not going to try and write knock-down objections to as yet unformulated justifications.

9.5 Caprice

To set out the correct decision-rule, Caprice, I need a new piece of terminology. Say ψ is *almost preferred* to ϕ according to P iff for all Pr in P , $E_{Pr}(\psi) \geq E_{Pr}(\phi)$. When no ambiguity results I omit the 'according to P '. Clearly whenever ψ is strictly preferred to ϕ it is almost preferred, but the converse is not true. Unlike strict preference, almost preference is not anti-symmetric. Bets ψ and ϕ can each be almost preferred to the other.

The core idea behind Caprice is that there should be as few restrictions on rational choice as possible apart from the rule that, whenever ψ is strictly preferred to ϕ it is irrational to choose ϕ over ψ . Unfortunately, as it is, this won't do, because it permits the following irrational course of action. Recall the earlier example where ϕ and ψ are almost indifferent, as are $\phi + \$5$ and ψ . If there were no rational restrictions on trade between almost indifferent bets then there would be no grounds for criticising the trader who first swaps $\phi + \$5$ for ψ and then swaps ψ for ϕ . Yet presumably it should be possible to subject this person to rational criticism.

I think the best thing to say about this case is that neither trade is itself irrational, but they are an irrational combination. In most decision-theories on the market this option is ruled out by stipulation; a set of trades is irrational iff one member of that set is irrational. There is, however, no reason to make such a restriction. Consider this analogy with belief. It seems plausible to say that it is reasonable to believe Oswald killed Kennedy and reasonable to believe he didn't, but it isn't reasonable to believe both that Oswald killed Kennedy and that he didn't. A set of beliefs, each reasonable on its own, might be unreasonable in combination. I am simply claiming we can say the same about decisions. A set of decisions, each reasonable on its own, might be unreasonable.

Because of this intuition, the Caprice rule must be expressed in terms of the reasonableness of sets of decisions. This can be applied easily to simple choices by looking at singleton sets. The notation $\#(\alpha, \beta) = \delta$ ($\delta \in \{\alpha, \beta\}$) means that δ is chosen (by the agent under consideration) in a pairwise choice between α and β . This is a different concept to the earlier $C(\alpha, \beta)$ notation in two respects. First, it is descriptive not normative. Given that I am usually discussing ideal agents this isn't as big a difference as it might normally seem. Secondly, $\#(\alpha, \beta)$ can be defined, even for rational agents, when $C(\alpha, \beta)$ is not. If α and β are almost indifferent, but when faced with the choice between them the agent chooses α , then $C(\alpha, \beta)$ is undefined (according to Caprice), but $\#(\alpha, \beta) = \alpha$.

Caprice: A set S of choices of the form $\#(\alpha_i, \beta_i) = \alpha_i$ ($i \in \{1, \dots, n, \dots\}$) is rationally permissible according to P iff there is some non-empty subset G of P such that for all i , α_i is almost preferred to β_i according to G .

Caprice is only defined in terms of pairwise choices. If α is chosen in a three-way choice between α , β and χ , we say $\#(\alpha, \beta) = \alpha$ and $\#(\alpha, \chi) = \alpha$. This can easily be extended to n -way choices. Hence a single n -way choice, with $n > 2$, can be regarded as a many-element set of pairwise choices.

Note two immediate consequences of this rule. First, when we are just considering a single choice between almost indifferent bets ϕ and ψ , either choice is acceptable. In trading terms, it is permissible but not compelling to trade ϕ for ψ . This is the motivation for calling the rule 'Caprice'. Secondly, any set of choices which leaves the trader with a position such that they would strictly prefer to be back where they started is not rationally permissible according to Caprice. Hence Caprice as specified captures the two important intuitive requirements on decision-rules.

I haven't yet specified how Caprice should be applied to choices between nodes of a decision-tree, because here there isn't much to say. In cases like that set out in Figure 9.1, the Capricious decision-maker can simply decide which end-point she wants to end up with, and follow the tree to that point. Provided her original n -way choice is permissible, every pair-wise choice she makes will be permissible. I showed above that the only way for the Conservative to avoid absurd decisions was to be closed-minded in the sense that she had to deliberately decide *not* to reflect at various stages in the tree about whether her initial strategy should be carried through. By comparison, the Capricious agent can be completely reflective.

There is one interesting special case of Caprice, which I'm adopting from Smets (1994). It isn't Smets's preferred approach for a couple of reasons, not least being that Smets advocates using the Dempster-Shafer updating rule, but the terminology and idea is largely his. An agent whose degrees of belief are vague over the set of probability functions in \mathcal{P} , whose 'representor' in van Fraassen's terms is \mathcal{P} , has \mathcal{P} as their *credal* probability function. They arbitrarily select an element P_i from \mathcal{P} to use for decision-making purposes; this is their *pignistic* probability. ('Pignistic' is from the Latin *pignus*, meaning to bet.) When making a choice between gambles they choose that gamble α such that $E_{P_i}(\alpha)$ is maximised. An agent who does this will never do anything wrong according to Caprice.

I noted at page 275 that any decision-rule would have to give up one of Arrow's constraints (1) through (4). Caprice gives up (2). It says that sometimes given the composition of \mathcal{P} we simply can't say which of two bets should be chosen. If this pignistic approach is followed, in a sense (2) is kept at the cost of (3). The pignistic probability function becomes the dictator in Arrow's sense. This might be an improvement; I leave it up to the reader to decide whether or not it is.

There is one odd result as a consequence of adopting Caprice. An agent is told (reliably) that there are red and black marbles in a box in front of them, and a marble is to be drawn from the box.

They are given the choice between three bets. α pays \$1 if a red marble is drawn, nothing otherwise, β pays a certain 45 cents, and χ pays \$1 if a black marble is drawn. Is it rationally permissible for the agent to choose β , again assuming constant marginal utility of money?

Levi (1974) writes as if it is obvious that choosing β is irrational. This is a cornerstone of the ‘impeccable’ analysis which leads to a dismissal of (4) but receives almost no justification. Jeffrey (1983) defines Bayesian approaches to decision-making so that choosing β is not Bayesian, but of course it isn’t an obvious truth that only Bayesian approaches are correct. Dempster (1988) claims that choosing β is permissible, and perhaps even compelling, though it appears he is motivated by the maximin rule, which I showed above is flawed.

I only bring this up to note that Caprice says it is not rational to choose β . To see this, assume we choose β . We will now show that G must be empty. Let p be the proposition that the marble to be drawn is red. Since β is almost preferred to α according to G , for every Pr in G it follows that $Pr(p) \leq 0.45$. However, since β is almost preferred to χ according to G , for every Pr in G it follows that $Pr(p) \geq 0.55$. There is no Pr satisfying each of these constraints, hence G is empty. It doesn’t however, appear at all intuitively compelling that it should be irrational to choose β . A defender of Caprice has to either explain away this intuition or, like Levi, simply deny that the intuition exists. The first of these choices is possible. One approach already noted is to say a choice of β reflects an irrational commitment to Maximin. Another is to say that it reflects a failure to internalise fully the assumption that the marginal utility of money is constant. I suspect that is what explains my intuition that β is an acceptable choice. I don’t think this raises a huge problem for the defender of Caprice – some questions are always going to be spoils to the victor – but it is a little disconcerting. If there is to be a strong attack on Caprice, I suspect it will be built around cases like this one.

9.6 Arguments For Caprice

Apart from the fact that it avoids the pitfalls of its more well-known rivals, there are two positive arguments for Caprice. Each of them is essentially the reverse of an argument I used against Levi. I'll call them the arguments from Arrow and Buridan.

The argument from Arrow notes that the four principles Arrow gave, (1) to (4) above, are inconsistent. Hence we must give up one of them. As there is strong intuitive support for Pareto, Non-Dictatorship and Independence of Irrelevant Alternatives, it seems the correct decision theory must give up what Arrow calls 'Collective Rationality', but what is perhaps better called Completeness in our context. There must be some choices about which our decision theory is silent. Since Caprice, unlike its popular rivals, satisfies this constraint, this is something in its favour. Of course this is not an argument against other incomplete rivals of Caprice. However, one strength of Caprice is that the class of decisions over which it is silent is quite a natural class. I doubt there could be a smaller class than this which is equally natural.

This leads to the argument from Buridan. Given the way I have set out the problem, when ϕ and ψ are almost indifferent, there is no reason to choose one over the other. The agent really is in the position of Buridan's ass. Of course like the ass the agent may be well advised to choose either ϕ or ψ over some less attractive alternatives. Unlike all its rivals, Caprice takes this conclusion seriously. If there is no reason to choose ϕ over ψ or *vice versa*, there really is no reason. It doesn't go and say this and then find a reason.

In particular, it must be really inexplicable why an agent chooses ϕ over ψ or *vice versa* in such cases. Should there be such a reason, it must be traceable to the beliefs and desires (or more generally partial beliefs and preferences) of the agent. The assumption of incomparability is just the assumption that those beliefs and desires don't determine a choice. Hence any decision theory must agree with Caprice's 'no explanation' conclusion. Given this, it is hard to see how the theory can differ from Caprice.

It might be thought that Caprice breaches this ‘no explanation’ rule in an important case. Say the expected value of ϕ is vague over [$\$30, \40], and that an agent has just sold a unit of ϕ for $\$32$. According to Caprice, if she now buys a unit for $\$38$, or indeed any price over $\$32$, she will have acted irrationally. Does this mean that either (i) the value of ϕ is now vague over merely [$\$30, \32] or (ii) the value is unchanged but she now has a reason for not buying ϕ for more than $\$32$? According to the objection, I have ruled out (i) and (ii), but I am committed to one of them.

The objection is in part correct, I have ruled out (i) and (ii). However, I am not committed to their disjunction. Were the agent to now buy ϕ for $\$38$, that would not of itself be an irrational act, however it would take her from having performed a set of rational acts to having performed a set of irrational ones. The only reason one would think this implies the last act is irrational is if one was wedded to the idea that a set of acts is irrational iff it includes an irrational act. By that principle, an agent can only move from a rational to an irrational set by performing an irrational act. However, that is a principle I gave reasons for rejecting in setting out Caprice.

Again the analogy with belief is instructive. If the agent believed yesterday that Oswald killed Kennedy, she can’t rationally believe today that Oswald didn’t kill Kennedy unless she ceases to believe that he did kill him. But, and here’s the difference, yesterday’s beliefs can be more easily undone than yesterday’s trades. If she could cease to have sold ϕ for $\$32$ yesterday, she can rationally buy it for $\$38$ today. Sometimes this will be possible (if the sale has a ‘cooling off’ period), but usually it will be just as fixed as the rest of the past. It is because she can change her beliefs, but not her trades, that we judge an agent’s trades diachronically, but her beliefs largely synchronically⁹. When we keep all this in mind, we won’t unduly focus on her last trade and judge it too harshly.

⁹ At least when she is moving between acceptable beliefs. If precision were mandatory, then less flexibility would be permitted.

9.7 Monte Hall Again

Now that I have the correct rule for updating vague degrees of belief (conditionalisation) and for decision-making with vague degrees of belief (Caprice) I can provide some advice to our contestant on the MHP. Recall that the contestant is, *ex hypothesi*, completely ignorant about which rule the host will use for showing her a door given that she has chosen the door which hides the car. Let p be the proposition that she has chosen this door. Her updated degree of belief in p should be vague over the interval $[0, 1/2]$. The bet ϕ pays whatever is behind the door she has chosen, and the bet ψ pays whatever is behind the other closed door less some small amount. Hence ϕ and ψ are almost indifferent, so she can rationally choose either door.

Note that if she had been committed to making decisions according to Levi's Rule or Conservatism she would be committed to keeping ϕ , that is not switching¹⁰. I have argued above that any of these options are errors, but if my arguments fail they provide an argument to the effect that switching is irrational. However, this only goes through if the representation of ignorance I have used is correct, as does the argument from conditionalisation and Caprice that it is not rationally compelling to switch.

And it seems this representation is a mistake. I was assuming that it was rational for the contestant's degree of belief in ab to be vague over $[0, 1/3]$. (' ab ' is the proposition that the prize

¹⁰ None of these approaches justify the rather odd conclusion of Moser and Mulder (1994) that it would be irrational to switch doors in a one-off Monty Hall game, but it would be irrational to stay if playing a repeat version of the game 100 times. To see why this is odd, consider the first of these games. By assumption it is better to not switch doors in the one-off game. So the cost of switching must be offset by gains in later games. But the games are totally independent, so what is this gain? Alternatively, and perhaps more controversially, run a backward induction argument showing that by Moser and Mulder's lights it is wrong to switch on the last game of the 100, hence it must be wrong to switch on the 99th, and so on. From their text it appears they have applied Dempster-Shafer updating rule (or something like it) to the single case. However, when considering the long-run case good common sense has overcome bad theory.

is behind the door she chose, a , and the host will open door b .) This means that it would be permissible to be disposed not to pay any arbitrarily small sum for a bet which pays an arbitrarily large sum if ab , and nothing otherwise. This, on reflection, seems unreasonable behaviour. It might be reasonable to refuse such a bet if it were offered (the offer being good evidence against ab), but the disposition seems unreasonable. I don't want to argue the contestant must have degree of belief precisely $1/6$ in ab , as is sometimes suggested. (See, for example, Horgan (1995).) Reason it seems would have our contestant be vague over an interval around $1/6$, but not all the way to 0. Since the possibility of the interval going so low is central to the argument that it is permissible to not change doors, that argument fails. It is, as many have suggested, unreasonable to stay with the original door.

Chapter 10

Keynes and Probability

10.1 Introduction

I have earlier described the theory defended in this thesis as a Keynesian theory of probability. In this chapter I want to defend that description. The purpose of this is twofold. First, there is some scholarly interest in accurately tracing the ancestry of ideas. Secondly, and more importantly, Keynes thought that his theory of probability led to distinctive theories in economics. I don't think he's right about this, for reasons I'll set out in the next chapter, but my argument turns on the theory of probability defended here being faithful to Keynes's account.

In the next section I'll set out the theory of probability Keynes advanced in his *Treatise on Probability* (1921a), and show how my theory captures the crucial elements of his. In some respects Keynes's theory is internally inconsistent: the detailed rendition does not capture the explicit motivating ideas. The theory here is closer to the spirit of Keynes's approach in these cases than its letter. Whether this is taken as charitable interpretation or friendly revision is probably not very important.

In section three I'll note how Keynes incorporated his philosophical insights into his economics. This will be elaborated in the next chapter, when I look at the details of part of his theory of unemployment. The most important element of this section is a theory of Keynesian 'uncertainty' which there is strong evidence to show he adopted.

There has been much debate in recent years over precisely what theory of probability Keynes was using when he wrote his *General Theory* (1936), and in particular how much the attacks of Ramsey and, perhaps, Wittgenstein had led him to abandon claims he had made in the earlier book. Much of this debate seems to assume that Keynes's early theory was such a tight package that it couldn't be amended without being lost. Hence we have arguments about whether Keynes

kept or abandoned his early theory. There is no doubt that he conceded some attacks were justified, but the evidence is consistent with his abandoning no more than a couple of relatively unimportant theses. In the last section I'll look at the claim he adopted a radically different theory of probability, a conventionalist theory. Again I argue the evidence is consistent with no change having occurred, and, a little more tentatively, I argue that it is inconsistent with his having adopted a conventionalist view. For convenience, in what follows I'll refer to the *Treatise on Probability* and *General Theory* as the *TP* and *GT* respectively¹.

10.2 Keynes's Pure Theory of Probability

The *TP* was Keynes's first major academic work. An early version of it was submitted as a fellowship dissertation in 1908, and after that was rejected a revised version was submitted in 1909. For various reasons it was not published for 12 years, with again significant revisions being made. Included in these revisions was a long discussion of statistical reasoning which arose from a public debate he had with Karl Pearson about the foundations of statistics. Although there is little evidence of it in his early books, Keynes started working theories about expectations into his economics as early as 1910 (Skidelsky 1983: 208; Keynes 1910: 46-7).

Keynes said hardly anything about probability for 10 years after his book was published. He didn't even make much of an effort to respond to scholars who contacted him about it, with more worldly concerns being pressing. In 1931, he made some passing comments in his review of Ramsey's

¹ Strictly speaking, it wouldn't hurt my story if Keynes had moved to a 'vague subjectivist Bayesian' position as in, say, Jeffrey (1983). This is precisely the position Runde (1994) has him adopt. I don't think a great deal in his economics depends on probability relations being objective, though everything turns on their being vague. Since my story would be hurt if Keynes had made the amendments suggested by others (e.g. Bateman (1996)) a small interpretative detour is justified.

book. Then in the *General Theory*, published in 1936, and the subsequent debate, he made some references to his theory of probability.

The basis of Keynes's theory is that probability is reasonable degree of belief, and this is a logical, non-numerical relation of 'arguments', which are just ordered pairs of propositions². The probability of an argument is what it is by virtue of logical laws alone, in just the way that an argument is either valid or invalid by virtue of logical laws alone. Indeed, Keynes thinks that the theory of valid arguments is just a special case of his theory of probable arguments, an argument being valid just in case it has probability one.

Not only are these probability values non-numerical, they are not always comparable. Probability is a non-linear as well as a non-numerical relation. In Keynes's terminology, the probability of an argument could be α , where α is a non-numerical degree of belief, comparable to some but not all numerical degrees of belief and to some but not all non-numerical degrees of belief. However, he says nothing about what type of entity α must be to have these properties. Because of this some later writers (eg T. Fine (1973), Runde (1994a)) have thought that Keynes's theory was best classified as a theory of comparative probability. In one sense it isn't, because it does make sense to talk about the probability of an argument being equal to some number, not just to compare the probabilities of different arguments. However, in a more important sense they are correct, because we know nothing about entities like α – which are the value of probability relations – except how they enter into comparisons. It seems that saying the probability of an argument is α is just shorthand for saying which comparisons between it and other probabilities are true.

Keynes might dispute this last analysis. He does develop a rather sophisticated calculus of these values, which impressed some rather distinguished logicians at the time. However, this calculus gives little insight into the nature of the values. As I remarked in chapter 6, when we derive

² By non-numerical I don't mean that it never takes numerical values, just that sometimes it does not.

$\alpha + \beta = \chi$ in a Keynesian system, where the Greek characters refer to the values of probability relations, there is no guarantee that '+' means what we mean by *plus*. In the absence of such a guarantee, or a more comprehensive analysis of what such derivations mean, it seems perfectly fair to classify Keynes's theory as comparative.

The other distinctive feature of Keynes's theory, and possibly the one which has aroused the most interest in the subsequent literature, is his concept of 'weight'. The probability of p given q tells us, in a sense, how the evidence q is balanced between p and $\neg p$. Keynes thought it might be important in some contexts, particularly decision making, to know not only how the evidence was distributed, but how much of it there was. Thus for every argument he postulated not only a probability, but also a 'weight', which is intended to be a measure of the comprehensiveness of the evidence. If this all sounds a bit vague, it is because the original is just as vague.

As the relevant evidence at our disposal increases, the magnitude of the probability of the argument may either decrease or increase, according as the new knowledge strengthens the unfavourable or the favourable evidence; but *something* seems to have increased in either case, – we have a more substantial basis upon which to rest our conclusion. I express this by saying that an accession of new evidence increases the *weight* of an argument. New evidence will sometimes decrease the probability of an argument, but it will always increase its 'weight' (*TP*: 77, italics in original).

Keynes modifies this slightly, saying strictly irrelevant evidence does not increase the weight of our argument. Finally Keynes introduced a new notation, which he thought essential to the progress of the subject. Given the lack of popularity of his notation, and the progress in the theory of probability since his time, his predictive skills seem to have failed here. However, for

expository purposes, it is sometimes convenient to adopt his notation. The essentials are to use $p / q = \alpha$ to mean the probability of p on evidence q is α , where α may, but need not, be a real number, and to use $V(p / q) = \beta$ for the weight of the argument from q to p .

The rest of this section aims to show how the theory of this dissertation captures the essential elements of Keynes's theory. Some parts of the translation are straightforward. Let \mathbf{P} be the set of all reasonable probability functions. Then the set P of values which $Pr(p | q)$ takes, for $Pr \in \mathbf{P}$ corresponds in a straightforward way to Keynes's non-numerical degrees. These sets are sometimes comparable with numbers. For real x , we say $P > x$ iff for all elements y of P , $y > x$. And this holds if we replace ' $>$ ' throughout with '=', '<', '≥' or '≤'. And sometimes the degrees cannot be compared with a real, as when none of these happen. Similarly the sets can sometimes be compared with each other, and sometimes not. Using Keynes's notation, $p / q > r / s$ iff for all $Pr \in \mathbf{P}$, $Pr(p | q) > Pr(r | s)$. Again this holds if we replace ' $>$ ' throughout with any other comparative. And since it is possible that no such comparative sentence is true, probabilities can be incomparable. Finally, we can now add probabilities in a way that makes it clear we are using the ordinary concept of addition. $p / q + r / s$ is the set $\{x: \exists Pr \in \mathbf{P} : Pr(p | q) + Pr(r | s) = x\}$. A similar definition applies to multiplication and division³.

So far the connection seems quite close. Since the membership of \mathbf{P} is non-contingent I almost capture Keynes's idea that the value of p / q can always be determined by logic. As I showed in chapter 1, there are good grounds for parting with Keynes on this point, but not for going all the way to subjectivism. A bigger departure comes with the translation of weight.

³ For some purposes it might be better to translate p / q as a function from \mathbf{P} to $[0, 1]$, which takes every Pr to $Pr(p | q)$. Then the talk of equality, inequality, incomparability, addition and multiplication is interpreted just as we ordinarily would for functions. For example where f and g are functions with a common domain $f > g$ iff for all x , $f(x) > g(x)$, and so on for the other definitions.

Keynes's explicit pronouncements about the weight of arguments are inconsistent with his notation and with the underlying idea behind his concept. I think both problems arise because Keynes says that even though $p / q = r / q$, $V(p / q)$ might be different from $V(r / q)$. The analysis of weight I will adopt denies this, and thus avoids the pitfalls that beset Keynes.

The notational flaw first. By representing the weight of the argument from q to p as $V(p / q)$, Keynes writes as if weight is not a two-argument function, with arguments p and q , but a one argument function, with the argument p / q . Recall that p / q is Keynes's notation for what I'd write $Pr(p | q)$. This confusion is evident at the top of page 79 of the *TP* where Keynes says "Let us represent the evidential weight of *the* argument, whose probability is p / q , by $V(p / q)$ " (my italics). Taken literally this seems to assume that there is only one argument with probability p / q , that is the one from q to p . Since this can't be right, the most natural reading of this sentence is to replace the italicised 'the' with 'any'. Alternatively, we should write the weight $V(p, q)$.

The problem is not, however, just a notational one. There is an argument that arguments of equal probability should have equal weight, something which Keynes denies. The problem arises because of disjunctive evidence.

One of the reasons Keynes wants to say that arguments of equal probability can have different weights is because he thinks that when $p / (q \wedge q_1 \wedge q_2) = p / q$, but not $p / (q \wedge q_1) = p / q$, then $V(p / q \wedge q_1 \wedge q_2) > V(p / q)$. This follows from his strict principle of irrelevance. Information q_3 is strictly irrelevant to the argument from q to p iff there is no conjunction $q_1 \wedge q_2$ such that $q_1 \wedge q_2$ is logically equivalent to q_3 and it isn't the case that $p / q \wedge q_1 = p / q$.⁴

⁴ This is to be distinguished from $p / q \wedge q_1 \neq p / q$. By the formal definitions of comparatives I've given, that would only be true if for all $Pr \in \mathbf{P}$, $Pr(p | q \wedge q_1) \neq Pr(p | q)$. The expression in the text is intended to be true just in case this holds for some Pr or other. If the theory of chapter 5 is correct, the two expressions have identical semantic content, but different conventional content. That, of course, is enough

This move to a strict principle of irrelevance appears to make it impossible for Keynes to ever say that a proposition is irrelevant to an argument. Let q_3 be any proposition such that $p / q \wedge q_3 = p / q$ and q_3 is in an intuitive sense irrelevant to the argument from q to p . For example, let p be the proposition that the die I am about to roll will land with six facing up, q be the background evidence, and q_3 be that it is now raining in New York. Now the extra evidence is divisible into two components $q_3 \vee p$ and $q_3 \vee \neg p$. The conjunction of these components is logically equivalent to q_3 , and each of them appear to be relevant pieces of information. Unless $q_3 / q = 1$, $p / q \wedge (q_3 \vee p) > p / q$. So it seems by Keynes's test, whether or not it is raining in New York is relevant to whether or not this die will land six. This information is relevant just because it is divisible into relevant components. Keynes notes this general move when discussing the principle of indifference on page 65 of the *TP*, but doesn't pick it up here. Carnap (1950: 420) seems to be the first to note that Keynes's strict concept of irrelevance is degenerate for the reason listed here.⁵

What is less frequently noticed in the Keynes literature is that the strong theory of irrelevance is central for Keynes's theory of weight. We are trying to find out whether or not p , and our current information is q . Assume that the strict concept of irrelevance is inapplicable when evaluating weight, so there is some information q_1 which is irrelevant, but which is divisible into relevant components. The weight of our belief in p (i.e. the weight of the argument from our evidence to p) will be unchanged if we find q_1 , as it is irrelevant. However, if we discover q_1 by discovering its components sequentially then, if weight increases upon the acquisition of relevant evidence,

to communicate the relevant point!

⁵ It seems quite amazing that no one picked this up before Carnap. Whitehead read the *TP* twice as one of Keynes's fellowship examiners, and Russell presumably read it closely before giving a strong endorsement of its formal work in Russell (1922). And while the other logicians with whom Keynes was in contact weren't in the same class as Russell and Whitehead they were still quite proficient. And Keynes notes that this decomposition of a proposition causes some difficulties for his rendition of the Principle of Indifference just five pages after he states the strict principle of irrelevance.

the weight of our belief will have risen. So the weight of our beliefs will depend not just on what our evidence is, but on the order in which we acquired it. This is absurd, so it must be that weight doesn't always increase upon acquiring relevant evidence, and may in fact decrease.

This argument reinforces the intuitive idea that when $p / q = r / s$ then $V(p / q) = V(r / s)$. The analysis of weight in my theory does this, as well as capturing I hope the spirit of Keynes's theory, although it departs radically from the letter of it⁶. Let P again be the set of values that $Pr(p | q)$ takes for $Pr \in P$. Then P will be an interval, and $V(p / q)$ is simply the one less length of that interval.⁷ As noted in the last chapter, the possibility of dilation means that we might be certain that on reception of the result of a trial (i.e. learning q_1 or learning $\neg q_1$) the weight of the argument from our evidence to p might go down.⁸ This is odd, but considerably less odd than the alternatives. It is a consequence of my analysis of weight that when we know the chance of p , the

⁶ There is a sense in which my departure from 'what Keynes really meant' isn't that great. If we adopt semantic externalism, then Keynes's references to 'weight' will refer to whatever it is in the world that best satisfies his descriptions of the role it plays. And that, I am confident, is what I have captured. Similar remarks will apply to 'probability' and other terms. Such an interpretation takes Keynes's intention for his referring terms to have referents more seriously than his particular description of those referents.

⁷ A similar suggestion for capturing this concept is made in Kyburg (1961: 225), though the interval in question has a somewhat different meaning in his theory.

⁸ This is also a consequence of the quite different analysis advanced in Runde (1990). He suggests we analyse weight as the ratio of known to unknown evidence, then immediately gives this an epistemic interpretation, so that weight becomes a ratio between how much we know and how much we know we don't know. Since new evidence might show that there is even more evidence we don't know, weight might go down with new evidence. I don't see, however, what the grounds are for the epistemic turn here; weight in Keynes seems to be objective enough that there could be a divergence between its value and our best estimate of it.

weight of our argument to p is maximised. As we'll soon see, this is a position which Keynes adopted when he applied his theory to economic matters in the 1930s.

10.3 Keynes's Applied Theory of Probability

Since the publication in 1983 of the first volume of Skidelsky's biography of Keynes there has been a surge in interest in Keynes's philosophy. A large motivation for this is the possibility that Keynes's philosophical beliefs play an important role in his economic theories. Like all modern economists, Keynes was aware of the role of expectations in determining the value of goods. Whenever we consider deferring consumption today in the hope of greater consumption tomorrow, our decision will depend crucially on our expectations of what tomorrow will be like; on our partial beliefs about tomorrow. The story is in fact a little more complicated. Since tomorrow we might decide to defer consumption again, it also depends on our expectations about two days time, and hence about three days time, and so on.

Here there is obviously a role for some theory of probability. However, many theories of probability fail to capture an important fact about this decision situation. We know more about tomorrow than about next year and more about the next year than the next decade. Keynes can represent this by saying that the weight of the argument from our evidence to propositions about tomorrow is greater than that of the argument from our evidence to propositions about next year. However, there is no way in which theories which say our degrees of belief ought be precise can do this. This applies whether the theories are subjectivist (like Ramsey) or logical (like Carnap). So Keynes has extra resources to bring to this problem, and he uses them.

It would be foolish, in forming our expectations, to attach great weight to matters which are very uncertain. (In footnote:) By "very uncertain" I do not mean the same thing as "very improbable". Cf. my *Treatise on Probability*, chap. 6, on "The Weight of Arguments". (GT: 148)

He elaborated on this in a famous passage from a 1937 article responding to critics of his theory.

By 'uncertain' knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty; nor is the prospect of a Victory bond being drawn. Or again, the expectation of life is only slightly uncertain. Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth holders in the social system in 1970. About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know. (1937a: 113-4)

Finally, in some 1938 he indicated the effect this uncertainty might have on decision theory. The context is a letter he wrote to Hugh Townshend, a former student, who appears to have been the first to inquire about the connection between the *TP* and *GT*.

One arrives presumably at the numerical estimates [of expected values] by some system of arranging alternative decisions in order of preference, some of which will provide a norm by being numerical. But that still leaves millions of cases over where one cannot even arrange an order of preference. When all is said and done, there is an arbitrary element in the situation. (1938a: 289)

O'Donnell (1991) claims there are obviously distinct senses of uncertainty underlying these three quotes. Call these the 1936, 1937 and 1938 senses, by reference to the year from which the quote

is drawn. I doubt that there are such distinct senses in play. There is a simple concept of uncertainty which can perform all the roles Keynes asks of it, and so the simplest explanation of the data is that Keynes was using that concept.

'Uncertainty' has been distinguished in economics from pure risk (such as we face at the roulette table) ever since Knight's *Risk, Uncertainty and Profit* (1921). However, Knight's distinction was based on a frequentist theory of probability, so it isn't overly important here. Say that p is uncertain for an agent with evidence q iff p / q is non-numerical⁹. More generally, say p is more uncertain the smaller the weight of p / q .¹⁰

Since non-numerical probabilities are usually associated with having little evidence, this captures the 1936 sense of uncertainty¹¹. The examples from the 1937 quote seem to match my definition, at least given some plausible assumptions about when the probability of p would be numerical. I also suspect we can translate 'calculable' from the last line as 'numerical'. Finally, non-numerical probabilities will easily lead to non-comparable goods. Say the probability of p is vague over $[x, y]$, with $x < z < y$. Then a p -bet will be incomparable with z cents. Further, incomparability (as compared to indifference) can't arise without this type of uncertainty.

As there is a simple sense of uncertainty which is definable in straightforward terms (indeed the definition uses no terminology not in Keynes) and which does the work needed for all three

⁹ Or in the functional terminology of footnote 3, if p / q isn't translated to a constant function. This suggestion first appears in Lawson (1985).

¹⁰ I.e. if x and y are the greatest and least values respectively such that $x \leq p / q \leq y$, the uncertainty of p is measured by $y - x$.

¹¹ There are counterexamples, as in Monte Hall type cases, but we can be very confident that Keynes was ignorant of these. And we can be confident that these don't threaten the claim in the text as to what is 'usual'.

quotes, I disagree with O'Donnell's claims that Keynes was using uncertainty differently every time. The main difference between our interpretations is that O'Donnell thinks the 1937 quote relied on the 'unknown probabilities' which are also introduced in the *TP*. We can all be ignorant of logical entailment relations. 'Snow is white' entails (at least on traditional definitions of entailment) Fermat's Last Theorem, but most of us were ignorant of that until very recently. Similarly, thought Keynes, we can be ignorant of logical probability relations. O'Donnell claims this is what is being referred to in the last line of the 1937 quote. He is forced to make the (empirical) claim that we are more often ignorant of these relations when the evidence is small, and he does this. The first problem with O'Donnell's position is that it is hard to see what the evidence for this could be. I am quite certain of the probability of 'My carpet is blue' given just 'Snow is white'; it is represented in my terminology by the interval $[0, 1]$. The second is that given the audience for this paper (it was printed in the *Quarterly Journal of Economics*) it would be unusual to base one's theory on the logical shortcomings of agents. I agree with O'Donnell that these shortcomings are hardly sufficient to say Keynes's agents are 'dim-witted', but I still think it would be an odd assumption for Keynes to make in that forum without more explicit acknowledgment. The third problem for O'Donnell's interpretation is that it seems to make uncertainty an 'on-off' property, rather than one admitting of degrees, as in my account. Yet Keynes's examples show he intended uncertainty here to come by degrees. Even if none of these arguments work, on general methodological grounds I prefer my interpretation because it is the simplest one that explains all the data.

There is a more important question on which I am in almost complete agreement with O'Donnell. In his (1989) he puts forward the most sustained defence of the position that Keynes's philosophical outlook was, in broad terms, unchanged throughout his life¹². This is despite the

¹² The different senses of uncertainty aren't a change of view on O'Donnell's theory, they are different facets of a multifaceted concept.

evidence that Keynes adopted a conventionalist epistemology (which I'll discuss in the next section) and that in his 1931 review of Ramsey's book he appears to yield to Ramsey on some crucial questions. I am in general agreement with O'Donnell on these points and there is little need here to repeat his arguments. There is however one point on which I do think Keynes abandoned the *TP* position, and this may explain Keynes's apparent yielding to Ramsey on some points.

In the review of Ramsey, Keynes notes how the scope of logic, which looked so broad under Russell and Whitehead, was gradually reduced by Ramsey and Wittgenstein until it looked like logic had no content at all. I think this is a sign that Keynes abandoned the position that his probability relations were logical. He didn't give up on their objectivity, he just now thought there were objective standards of rationality that were extra-logical. These objective standards, as opposed to logic, would provide the foundations of probability. Without going for a line-by-line exegesis, I'll just note that Keynes's text supports the claim that it was on this point that Keynes yielded to Ramsey.

That much seems quite well grounded; what follows is rather speculative. A year before Keynes's review of Ramsey appeared, Gödel announced his incompleteness theorems. These were widely taken to cast doubt on Russell and Whitehead's program of incorporating all of maths into logic. If Keynes knew of this result and that reaction (which is at least possible, but is hardly certain) this might have convinced him there were objective epistemic norms (i.e. conformity with mathematical reasoning) which were not strictly logical, and this would have made his abandonment of the claim his probability relations were logical less dramatic. There is no reference to Gödel in Keynes's review, which might look like evidence against my speculation that he was influenced by Gödel on this point, but on the other hand there is no reference in it to anyone at all outside Cambridge!

This isn't intended as a criticism of theorists who say that Keynes's theory of probability was little changed over his life. I suspect that, from Keynes's point of view, this was a little change. Another little change he may have made was dropping the view that his probability relations were non-natural¹³. The difficulty facing all sides in this debate is that Skidelsky's summary of Keynes's later view seems unimpeachable: "[After 1921] Keynes no longer thought about probability, though probability permeated his thought." (Skidelsky 1992: 73) If Skidelsky's right, then the correct approach is what I've adopted. The only role for probability in Keynes's later thought is in application.

10.4 Keynes and Conventions

Many authors argue that Keynes adopted a conventionalist, intersubjectivist theory of probability. See, for example, Bateman (1996), Davis (1994), and Gillies (1988). I'll focus on Bateman's account, as it's the most comprehensive and most recent. It isn't always obvious what is meant by such a theory. In particular, it isn't clear whether it is meant to be an empirical or a normative theory; whether Keynes is apparently claiming that we ought set our partial belief by convention or that we in general do. Since the empirical theory would be consistent with his objectivist norms, and stress is on the change in his views, I conclude that the claim is that this is a new normative view. According to this view, being reasonable is analysed as conforming to conventions; this contrasts with the *TP*'s position that being reasonable involves just making reasonable judgements in particular cases. This is not a very standard epistemological position, but something similar is often endorsed in ethics. In ethical language, it is a move from an act-based to a rule-based epistemology. Bateman marshals the evidence that Keynes moves from an objectivist to a

¹³ I wonder if he might have been convinced they were not by Ramsey's analysis of theoretical terms (Ramsey 1929a). This could be pressed into giving a naturalist reading of apparently non-natural terms, but I rather doubt Keynes would have seen this, or even considered it an interesting object of study compared to his other projects at the time.

conventionalist position in ethics as evidence for this epistemological shift, but this doesn't seem of overwhelming significance¹⁴.

Here's the closest Bateman gets to a definition of what he means by an intersubjective theory of probability.

When probabilities are formed according to group norms, they are referred to as intersubjective probabilities ... I take it to be the case that in a world of subjective probabilities some individuals will form their own estimates and others will form them on the basis of group norms (50n).

This makes it look very much like an empirical theory, since it refers to how people actually form beliefs, not how they ought. So his intersubjectivism looks perfectly consistent with Keynes's objectivism. I am completely baffled by the 'world of subjective probabilities'. I wonder what such a world looks like, and how it compares to our world of tables, chairs and stock markets?

Fortunately there is a theory which does the work Bateman needs. Ayer (1936) rejects orthodox subjectivism about probability on the grounds that it doesn't allow people to have mistaken probabilistic beliefs. But he can't admit Keynesian probability relations into his sparse ontology. The solution he adopts is to define probability as degree of rational belief, but with this caveat.

¹⁴ If Keynes had adopted a framework which implied a tight connection between epistemological and ethical norms, such as utilitarianism, this would be important, since he couldn't change ethics and keep his epistemology. But such frameworks aren't compulsory, and, given the vehemence with which Keynes denounced utilitarianism (Keynes 1938b, 445), it seems he didn't adopt one. Further, there is substantial evidence that, by the time the *TP* was completed, Keynes had already adopted a rule-based ethics as in effect a response to the war. See Keynes (1921b: 298) for the most compelling such evidence. Hence there was no connection in Keynes's mind between ethics and epistemology on this point.

Here we may repeat that the rationality of a belief is defined, not by reference to any absolute standard, but by reference to part of our own actual practice (Ayer 1936: 101).

The ‘our’ is a bit ambiguous; interpreting it to refer to the community doesn’t do violence to the text, though it is just as plausible that it refers to a particular agent. The ‘part of our practice’ referred to is just our general rules for belief formation. These aren’t justified by an absolute standard; they are justified by the fact they are our rules, and presumably by their generality. Given Bateman’s views about metaphysics¹⁵, it seems quite reasonable to suppose he’d follow Ayer on this point.

The evidence that Keynes adopted such a position is usually taken to be some passages from the *GT* and the 1937 *QJE* paper I quoted above. Here’s the key points from the two quotes Bateman uses to support his view.

In practice we have agreed to fall back on what is, in truth, a *convention*. The essence of this convention – though it does not, of course, work out quite so simply – lies in assuming that the existing state of affairs will continue indefinitely, except in so far as we have specific reasons for expecting a change (*GT*: 152).

How do we manage in such circumstances to behave in a manner which saves out faces as rational, economic men? We have devised for the purposes a variety of techniques, of which much the most important are the three following: ...

¹⁵ He thinks that ‘analytic philosophy is the last place one would expect to find metaphysics’ (Bateman 1996: 39)

(3) Knowing that our own individual judgement is worthless, we endeavour to fall back on the judgement of the rest of the world which is perhaps better informed. That is, we endeavour to conform with the behaviour of the majority or the average. The psychology of a society of individuals each of whom is endeavouring to copy the others leads to what we may strictly term a *conventional* judgement (Keynes 1937a: 115).

There are two problems with using this evidence the way Bateman does. The first is the old one that they seem expressly directed to empirical questions, though in economic writings with their assumptions of rationality such appearances can be deceptive. The more important one is that Keynes is attempting to answer a very specific question with these passages; in ignorance of the question, we can easily misinterpret the answer.

How much ought one pay for a share in company X? Well, if one intends to hold the share come what may, all that matters is the expected prospective yield of X's shares, appropriately discounted, as compared to the potential yield of that money in other uses. But as Keynes repeatedly stresses (*GT*: 149; Keynes 1937a: 113-4) we have no basis for forming such expectations. Were this the only reason for investing then purely commercial investment may never happen.

There is another motivation for investment, one that avoids this problem. We might buy a share in X today on the hope that we will sell it next week (or next month or perhaps next year) for more than we paid. To judge whether such a purchase will be profitable, we need a theory about how the price next week will be determined. Presumably those buyers and sellers will be making much the same evaluations that we are. That is, they'll be thinking about how much other people think X is worth.

We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees (*GT*: 156).

There is simply no solution to this, except to fall back on convention. That is, we are forced into a conventionalist theory of value, at least of investment goods. But this doesn't mean that we have a conventionalist epistemology. On the contrary, it means that our ordinary (objectivist) empiricism is unimpeded. For the question that Keynes has us solve by reference to convention is, what is the value of X? This is equivalent to, what will be value of X be, or again, to what are the conventional beliefs about X's value? A question about the state of conventions needs answering; and, as good empiricists, we answer it by observing conventions.

An analogy may help here. Here's something that Hempel believed: to gain rational beliefs about the colour of ravens, one has to look at some birds. Did this mean he had an ornithological epistemology? No; he had an empiricist epistemology which when applied to a question about ravens issued the directive: Observe ravens! Similarly Keynes's belief that to answer questions about value, i.e. about conventions, one has to look at conventions, does not imply a conventionalist epistemology. It just means he has an empiricist epistemology which when applied to a question about conventions issues the directive: Observe conventions!

There might be another motivation for using conventions, again consistent with Keynes's objectivist empiricism. Sometimes we may have not made enough observations, or may not have the mental power to convert these to a theory. So we'll piggyback on someone else's observations or mental powers. This seems to be what's going on in the quote from Keynes (1937a). Or even better, we'll piggyback on everyone's work, the conventions. To see how this is consistent with an objectivist epistemology (if it isn't already obvious) consider another analogy.

What is the best way to work out the derivative of a certain function? Unless your memory of high-school calculus is clear, the simplest solution will be to consult an authority. Let's assume for the sake of argument that the easiest authorities to consult are maths texts. It seems like the rational thing to do is to act as if the method advanced by the maths texts is the correct method. Does this mean that you have adopted some kind of authoritarian metaphysics of mathematics, where what it is for something to be correct is for it to be asserted by an authority? Not at all. It is assumed that what the textbook says is correct, but the authoritarian has to make the extra claim that the answer is correct *because* it is in the textbook. This is false; that answer is in the textbook because it is correct. In sum, the authoritarian gets the direction of fit wrong.

Similarly, in the 'piggyback' cases, the intersubjectivist gets the direction of fit wrong. In some circumstances, I might assume that if p is 'average opinion', then it is reasonable to believe p . But I wouldn't say it is reasonable to believe p because p is average opinion, as the intersubjectivist does; rather, I say that assuming p is average opinion because it is reasonable to believe p .

The evidence so far suggests Keynes's statements are consistent with his denying intersubjectivism. I might be able to go further and show they are inconsistent with his adopting that theory. After the quote on *GT* page 152 he spends the next page or so defending the use of conventions here. The defence is, in part, that decisions made in accord with conventions are reversible in the near future, so they won't lead to great loss. If he really were an intersubjectivist, the use of conventions would either not need defending, or could be defended by general philosophical principles. Secondly, there is this quote which in context seems inconsistent with adopting a conventionalist view.

For it is not sensible to pay 25 for an investment which you believe will yield 30, if you also believe that the market will value it at 20 three months hence (*GT*: 155).¹⁶

The context is that he is discussing why reasonable professional investors base their valuations on convention rather than on long-term expectation. Hence the ‘you’ in the quote is assumed to be reasonable. Hence it is reasonable, Keynes thinks, to believe that an investment will yield 30, and that conventional wisdom is that it will yield much less. But if all reasonable beliefs were formed by accordance with conventional wisdom, this would be inconsistent. Hence Keynes cannot have adopted a conventionalist epistemology.

¹⁶ This might be considered odd if we take a market-orientated view of ‘yield’. We could plausibly just view the yield of an investment as the value the market will place on it three months hence. Indeed the argument in Chapter 12 of the *GT*, from where the quote is drawn, would support such a view. However, Keynes has made it clear from the context that he intends ‘yield’ to refer to the real returns of the investment over its lifetime, a question on which a sensible investor can disagree with the market.

Chapter 11

The Economic Consequences of Uncertainty

11.1 Uncertainty, Investment and Unemployment

“Employment was a problem because investment was; and investment was problematic because of the uncertainty of its return.” (Shapiro 1997: 83)

This is a nice two-line summary of an important part of Keynes’s theory of unemployment. A large part of the revival of interest in Keynes’s philosophy of uncertainty in recent times, as represented by its central place in five anthologies produced about Keynes this decade, is because of its connection to his radical views on unemployment. In this chapter I will investigate the idea that the distinctive features of Keynes’s theory of probability, in particular his allowance that degrees of belief be non-numerical, have economic implications. My main conclusion is negative; there is no distinctive philosophical theory here which is (a) correct and (b) has interesting economic consequences.

It is worth briefly going over the broader economic theory into which Keynes’s views about uncertainty are meant to fit. Unemployment, or at least involuntary unemployment¹, is a monetary phenomenon for Keynes. The two crucial premises, which are well beyond the scope of this dissertation, are: (1) that demand for money is not, even indirectly, demand for labour and (2) that decreasing the price of labour will not, *ceteris paribus*, increase demand for labour. Given these premises, a shift in demand from investment or consumption goods to money will lead to a

¹ To be distinguished not just from voluntary unemployment, but from ‘frictional’ unemployment resulting from the delay in matching demand and supply of labour. It has been argued in recent times that there is no involuntary unemployment in Western economies in the sense Keynes specifies. Given the breadth of the frictional and voluntary categories (‘voluntary’ unemployment includes unemployment caused by voluntary decisions of government to, say, prohibit some kinds of industry) this claim isn’t quite as absurd as it first appears.

decrease in demand for labour, and this decrease cannot be rectified by a change in price. Both premises are *very* controversial, and it is a little controversial whether this is the correct way to represent the argument. But these questions I set aside for a different project. Here the task is to see whether uncertainty might cause an increase in the demand for money, or equivalently a decrease in the demand for investment goods, and hence by these premises an increase in unemployment.

It should be stressed that the kind of uncertainty in question here is epistemic. Some writers, such as Paul Davidson (1994) and Rogers and Rymes (1997) have written as if some kind of aleatory uncertainty is needed for the argument to go through. “By an uncertain world, we mean one in which the knowledge which we must have in order to know fully, if stochastically, the consequences of our acts is in principle unattainable” (Rogers and Rymes 1997: 304). My quibble is with the ‘in principle’. For me, an uncertain world is one in which some rational agents are sometimes in a situation in which it is reasonable for them to have non-precise degrees of belief about propositions whose truth is relevant to their economic decisions. This imprecision may be ‘in principle’ removable if, for example, the chance of all propositions is precise², but that doesn’t make the situation any more certain.

² Is this true of our world? I hesitantly guess *yes*, just because the ways in which it could be false seem so implausible. Here’s the two best possibilities for constructing propositions with non-numerical chances. Perhaps the laws of chance only apply to ‘basic’ events, satisfy an additivity principle no stronger than denumerable additivity, and some propositions about which we are interested, say that the Yankees will win the next World Series, are really disjunctions of non-denumerably many basic events. In this case the proposition will behave like a non-measurable set, and have ‘inner’ and ‘outer’ chances, but no precise chance. Each of the premises here seems questionable. Second possibility. Perhaps for some systems the relevant laws of nature are not ‘dynamic’ laws, but only ‘comparative static’ laws. So it is a law that once a system is disturbed that it will revert, quickly, to an equilibrium state. However, when there are multiple equilibria there are no laws about which one will be actualised, and in practice there are no stable

Davidson's discussion suggests that there is a reason for resisting this characterisation of uncertainty. The problem he sees is that in a world with resolvable uncertainty the competitive pressures of the market will, in the long-run, select those agents whose degrees of belief match the chances³. There are three problems with this argument. First, even if it works and in the long-run uncertainty is resolved, it isn't clear why this is interesting, or at least why it ceases to be interesting to examine what happens in the short-run. As Keynes stressed, developing new variants on "When the storm is over the seas will be flat" is not much of a theoretical challenge. Secondly, when it is true, and known, that the costs of resolving uncertainty in this way will be prohibitive, that a resolution is in principle possible seems somewhat irrelevant. Imagine trying to determine who to bet on for the next baseball season by working out the objective chance of each team winning! Thirdly, and perhaps more controversially, it isn't obvious why in the long-run those agents whose degrees of belief match the chances will be selected. After all, the agents who are selected will be those whose predictions turn out to be correct. So lucky guesses will confer greater survival advantage than beliefs developed in accord with chances. Maybe the game goes for so long that luck will run out; maybe there are so many agents starting off who are just guessing that some of them will outlast those who know the chances. *A priori* it seems hard to say which is true. So if numerical chances for all events exist, it isn't clear that even in the long-run market participants will know what they are. Hence I feel confident in restricting my attention to epistemic uncertainties.

frequencies of reversion patterns. If markets, or ecologies, or complex machines are merely governed by such laws then again we may have events with imprecise chances. But again this seems just too mysterious to be real. Whether or not either of these speculations is correct does seem irrelevant to whether an investor in the copper market is really in an uncertain situation.

³ Davidson doesn't really run this argument, but is concerned that his classical opponents will have recourse to it if he allows uncertainty to be resolvable in principle.

11.2 Two Consequences of Uncertainty

Uncertainty impacts on demand for investments and money in two related ways. These ways are not distinguished in the literature, but it will help my exposition to have the distinction clear. First, uncertainty may reduce demand for investment *directly* by making a person who would otherwise be tempted to invest more cautious and hence reluctant to invest. Secondly, if this direct impact is widespread enough, it will increase the demand for money, and hence its price. But the price of money is just the market rate of interest. And the return that an investment must be expected to make before anyone – even an investor not encumbered by uncertainty – will make it is the rate of interest. I will call this the indirect impact of uncertainty on investment.

This needs explaining in some more detail. Keynes takes the amount of consumption as a given (*GT*: 245). Or more precisely, for any period he takes the amount of available resources that will *not* be allocated to consumption as a given. There are three possible uses for these resources: they can be invested, they can be saved as bonds or loans, or they can be hoarded as money. There are many different types of investment, but Keynes assumes that any agent will already have made her judgement as to which is the best of these, so we need only consider that one. There will also be many different length bonds which the agent can hold. So as to simplify the discussion, Keynes proposes just treating these two at a time, with the shorter length bond called ‘money’ and the longer length loan called ‘debts’ (*GT*: 167n). Hence the rate of interest is the difference between the expected return of the shorter bond over the life of the longer bond and the return of the longer bond. So the rate of interest that we are interested in need not be positive, and when the two bond lengths are short will usually be zero. It is usually presumed in discussions that the rate is positive, and I’ll generally follow in making that assumption⁴. Now, Keynes presumes that an agent will only allocate resources to investment if investment looks to be at least as worthwhile as holding

⁴ The empirical evidence seems broadly consistent with the theory that the rate of interest on an n -year bond is simply determined by the expected return on bank money over that time. On that theory all the rates with which we are concerned will be zero.

money, and at least as worthwhile as holding debts. In other words, he makes the standard reduction of n -way choice to a set of 2-way choices⁵.

Now, usually, if someone is of a mind to invest, they will not favour holding money over holding debts. The only motivation for the latter, given positive interest rates, could be a desire to have accessible command over purchasing power, and investment foregoes that command. So, in practice we only need look at two of the three possible pairwise choices here. I will ignore for most purposes the choice between investing and holding money, and only look at the money-debt choice and the debt-investment trade-off.

Holding a debt provides a relatively secure return in terms of money. Relatively secure because there is the possibility of default. In practice this means that there isn't a sharp distinction between debts and investments, rather a continuum with say government bonds at one extreme and long-term derivatives at the other. Some activities which have the formal structure of 'debts', like say provision of venture capital, will be closer to the investment end of the continuum. Unlike debts then, investments as a rule do not have a secure return in terms of money. In most cases they do not even have a precise expected return (*GT*: 149; Keynes 1937a: 113).

Now Keynes does not presume that this means that people never invest unless the expected return on the investment is greater than the expected (indeed, known) return on debts. He says explicitly that were this true then 'there might not be much investment'. Instead he says that investment under uncertainty depends on 'confidence' (*GT*: 150). So the following looks compatible with his position. Let the expected value of investing a certain sum be $[\alpha, \beta]$, and the expected value of buying a debt with that money be χ . Then the agent will invest iff $(1 - \rho)\alpha + \rho\beta \geq \chi$, where

⁵ Standard, but I bring it up because the modern theorist who's decision theory is closest to the one Keynes seems to adopt, Levi, explicitly rejects it.

$\rho \in [0, 1]$ measures the ‘state of confidence’⁶. Now when a crisis erupts, ρ will go to 0, and investment will dry up. In such cases the decision theory is similar to the one advanced by Levi. In general this is exactly the theory advocated by Strat (1990) and Jaffray (1994). Since we’re interested in a theory of unemployment, we’re primarily interested in the cases where ρ is quite low, in which cases we can say uncertainty is reducing investment.

That last statement might seem dubious at face value. In part what I mean by it is this. When ρ is low, the value of a set of bets will in general be *more* than the sum of the value of the bets taken separately. Because individual investors are fearful of exposure to uncertainty, which is presumably what ρ being low means, sets of investments which if undertaken collectively would be profitable (and everyone agrees that they would) will not be undertaken individually. This not only shows the important role uncertainty plays, but it suggests a reason that theorists have thought government intervention might be appropriate in times of crisis. Alternatively, if ρ is low then the value of an investment, how much we will be prepared to pay for it, will probably be lower than our best estimate of its expected return, assuming the latter to be near $(\alpha + \beta)/2$.

I’ll focus more closely on the ‘indirect’ effects of uncertainty in section 4. The central idea is that the rate of interest, being the price of money, is completely determined in the market for money. But this market itself has some rather strange properties. After all, money is barren, and it can generally be traded for something which is not barren. So, as Keynes puts it, why would anyone ‘outside a lunatic asylum’, want it? Why wouldn’t the demand for money drop to zero as soon as the rate of interest is positive?

⁶ In case the reader fears I am being absurdly formal with an essentially informal idea, Keynes had such a variable, there described as measuring the ‘state of the news’, in early drafts, but it didn’t survive to the final stage. So my proposal isn’t a million miles from what Keynes intended merely by virtue of being algebraic.

Because, partly on reasonable and partly on instinctive grounds, our desire to hold money as a store of wealth is a barometer of the degree of our distrust of our own calculations and conventions concerning the future. Even though this feeling about money is itself conventional or instinctive, it operates, so to speak, at a deeper level of our motivation. It takes charge at the moments when the higher, more precarious conventions have weakened. The possession of actual money lulls our disquietude; and the premium which we require to make us part with money is the measure of the degree of our disquietude (Keynes 1937a: 116).

So more uncertainty means more demand for money means higher interest rates. The rest of the story is standard. Even the confident agent will be disinclined to invest once the rate of interest rises. Using the little decision theory outlined above, more uncertainty means the gap between α and β grows, which if ρ is low will tend to reduce $(1 - \rho)\alpha + \rho\beta$, the ‘certainty equivalent’ of the expectation of the investment’s worth⁷. On the other hand, uncertainty on the part of the community will tend, for similar reasons, to increase χ . Either way, investment suffers, and hence so does employment.

11.3 Uncertainty and Money

There is something very odd, from an economic point of view, about all that we have done so far. Let’s take a simple example where concerns about uncertainty are supposed to lead to an agent holding money rather than investments. Say shares in Company X are currently priced at \$1, and for simplicity assume that X does not pay dividends, so the expected return on holding X’s shares

⁷ Keynes suggests at *GT*: 24n that for mathematical purposes we will have to reduce expectations to ‘certainty equivalents’, and seemed to hold this belief throughout the debate on the book.

is just the expectation of X 's share price at the end of the period⁸. Assume this expected return is vague over $[\$0.8, \$1.5]$. Then any agent such that $\rho < 2/7$ will prefer to hold money rather than dip into the market. In particular the agent with maximal uncertainty aversion, in a non-technical sense of aversion, who sets ρ at zero will hold money. But there is a sense in which this is exposing oneself to more uncertainty rather than less. The expected value of $\$1$, in terms of X 's shares, at the end of the period is *ex hypothesi* vague over $[2/3, 5/4]$. If we hold a share we are certain to have something that will be worth one share, however if we hold a dollar our expectation of its value can't be made more precise than being more than two-thirds of a share and less than five-quarters of a share. Surely if we dislike uncertainty we'll hold the share. Well, at least we deserve an explanation as to why one kind of uncertainty is given such a central place and other kinds are completely ignored.

Keynes has an explanation, and in a sense it's philosophically rigorous. He argues, or perhaps assumes, essentialism about money. Indeed the title of chapter 17 of the *GT* is 'The Essential Properties of Interest and Money'. These essential properties are entirely functional. As Hicks puts it, "Money is defined by its functions ... money is what money does" (Hicks 1967: 1). The sense in which this is rigorous is that the assumption is only that agents try to minimise uncertainty relative to whatever is money in their world, whatever plays the functional role of money. Hence the solution to the problem posed in the previous paragraph is guaranteed to be non-arbitrary⁹.

As a first approximation, we can say the functional role money plays is that it is a medium of exchange. Keynes doesn't think this is quite the essential property; rather he says that money is

⁸ I'm also abstracting away from tax considerations.

⁹ The essentialism here is, I think, different to the type of essentialism that has become philosophically popular since Kripke (1972). In particular we don't say here that money is whatever plays a certain functional role in the actual world. Rather we say that money is what plays a certain functional role in the world under consideration. It is not just *a priori*, but necessary that money has certain functional properties.

essentially ‘liquid’, and perceived to be liquid. This means that if we hold money we are in a position to discharge obligations and make new purchases as they seem appropriate with greatest convenience and least cost. Even this isn’t what’s given as the official essential property of money. To make the proof that demand for money is not demand for labour easier Keynes takes the essential properties of money to be its negligible elasticities of production and substitution¹⁰. But he makes clear that these are important because of their close connection to liquidity (*GT*: 241). Indeed, when he comes to define a non-monetary economy, he simply defines it as one where there is no good such that the benefits it confers via its liquidity, its ‘liquidity premium’ exceeds the carrying costs of the good. So the properties of having a negligible elasticity of production and substitution seem necessary but insufficient for something to be money.

That money is liquid will explain why uncertainty in terms of it is important. At the end of the day, the point of holding investments, bonds or money is not to maximise the return in terms of such units; it is to be used somehow for consumption. Hence we prefer, *ceteris paribus*, to store wealth in ways which can be easily exchanged for consumption goods as and when required. Further, we may be about to come across more information about productive uses for our wealth, and if we do we would prefer to have the least inconvenience about changing how we use wealth. Money is going to be the best store of wealth for each of these purposes. The strength of these preferences is the liquidity premium which attaches to money.

¹⁰ As all of Keynes’s units are defined in terms of labour, elasticity of production is defined to be “the response of the quantity of labour applied to producing it to a rise in the quantity of labour which a unit of it will command” (*GT*: 230). Put simply, elasticity of production measures the tendency of supply to rise in response to higher prices. Elasticity of substitution measures the tendency of consumers to choose goods other than the one in question as its price rises. Again, Keynes defines the price of money in terms of wage-units.

So Keynes's story here is essentially a 'missing markets' story. If there were markets for every kind of transaction there would be no liquidity premium attaching to money, and hence no reason to be averse to uncertainty in terms of money returns as opposed to uncertainty in terms of X's shares returns. It is worthwhile to note here a methodological difference between decision theorists and economists. In decision theory it is common to specify what choices an agent does have. These will usually be finite, or at least simply specified¹¹. In economics it is more common to specify what choices an agent does not have, which markets are 'missing'. In a sense the difference is purely cosmetic, but it can change the way problems are looked at. Since Keynes requires here some markets to be missing, it might be worth investigating what happens here from the more restrictive framework ordinarily applied in decision theory.

In some decision-theoretic contexts we can have a preference for liquidity even when we are completely certain about what our choices are and what their outcomes will be. Say we are in a game where the object is to maximise our money over 2 days. We start with \$100. On day 1 we have a choice of buying a ticket which will pay \$200 at the end of day 2, and is non-transferable, or doing nothing. On day 2, if we still have our \$100, we can buy with it a voucher which pays \$300 at the end of day 2, or doing nothing. Obviously the best strategy is to do nothing on day 1, and buy the voucher on day 2. The point is just that money here has enough of a liquidity premium on day 1 that we are prepared to hold it and earn no interest for that day rather than buy the ticket (or two day bond) which will earn interest. So uncertainty is not a necessary condition for liquidity premiums to exist. On the other hand perhaps it is necessary for liquidity premiums to exist in a world something like ours, where agents neither have all the choices they would have in a perfect market, nor as few as in this simple game. If we added a market for tickets and vouchers to our simple game the prices would be fixed so that money would lose its liquidity premium. Keynes

¹¹ An agent's choice may range over a continuum, but the description of the type of choices open will usually be quite brief.

suggests something like this is true for the worlds he is considering: “*uncertainty* as to the future course of the rate of interest is the sole intelligible explanation of the type of liquidity preference [under consideration]” (*GT*: 201). However, here he merely means lack of certainty; there is no proof that if every agent had precise credences liquidity preference ought disappear. So it looks like uncertainty in the sense I’m interested in, vague reasonable beliefs, does no theoretical work. Perhaps this is a bit quick, as the little game I considered is so far from a real-life situation. So I will look more closely at the effects uncertainty is supposed to have. Since it has received the bulk of the theoretical attention, I start with the indirect effects of uncertainty.

11.4 Uncertainty and Liquidity Preference

As noted above, Keynes thinks the question of why money is demanded at all, why we don’t all move from holding money into holding debts as soon as the rate of interest goes positive, needs answering. And he thinks the answer here will be particularly relevant to theories about the rate of interest. If the market in general is at equilibrium then the market in trades between any two goods must also be in equilibrium; in particular it cannot be that there are people holding money who would be prepared to buy debts at the current interest rate. So if the equilibrium interest rate is positive, there must be some people who would prefer to hold money than hold debts. This fact Keynes takes to be central to the correct theory of the rate of interest¹². Hence to determine what the rate of interest will be, and what will cause it to change, I need to determine what causes a demand for money.

¹² Saying why precisely Keynes takes equilibrium in this market, rather than say the market between debts and investments, to be of central importance would take us too far from our topic. If that market is taken to be central then we get a theory of the rate of interest in terms of the productivity of investment. In general Keynes’s reasoning here is that there are other markets which can adjust to bring equilibrium to the debt-investment market whatever the rate of interest, but this is not so for the debt-money market, but we need to say more to make this a full response.

Keynes distinguishes four motives for holding money (*GT*: Ch. 13; Keynes 1937b: 215-223). Two of these, the transactions motive and the finance motive, need not detain us. They just relate to the need to make payments in money and on time. The third, the speculative motive, is often linked to uncertainty, and indeed Keynes does so (*GT*: 201). But ‘uncertainty’ here is just used to mean absence of certainty, that is the existence of risk, which is not how I’m using ‘uncertainty’. The idea behind the speculative motive is that if an agent believes the rate of interest is about to rise (the value of debts is about to fall) it is better to hold money now and buy the debts after the change. This holding is described as speculative because it is just a gamble on a change in the market. Now if all players in the market were certain as to future movements of the rate of interest, there could be no speculative motive, so the speculative motive requires the absence of certainty, but this is not particularly interesting. How a roulette wheel will land is not certain, but we wouldn’t say it is subject to uncertainty in this sense. As Runde (1994b) points out, an agent who is certain as to future movements in interest rates may still hold money for speculative reasons, as long as other agents who are not so certain have made mistaken judgments. The fourth motive will hold most of my attention. Keynes argues that we may hold money for purely precautionary reasons.

To provide for contingencies requiring sudden expenditure and for unforeseen opportunities of advantageous purchases, and also to hold an asset of which the value is fixed in terms of money to meet a subsequent liability fixed in terms of money, are further motives for holding cash (*GT*: 196).

I should clarify what is meant by ‘unforeseen’ here. It doesn’t include all purchases or contingencies such that it was not foreseen that they would (probably) happen. Rather, it is limited to those purchases or contingencies such that it was not foreseen that something of that vague type would happen, or at least that not as many things of that type would happen as actually did. For

example, if some part on an old car breaks down and needs replacing this would be classed as ‘forseen’ not because it could be predicted that particular part would break down, but rather that some part or other would break down. Although Keynes is not explicit on this point, he makes just this distinction when distinguishing between voluntary and involuntary losses arising from the use of capital equipment (*GT*: 56) and it seems reasonable to assume he’d make a similar distinction here.

There are a number of attempts in the literature to explain why uncertainty leads via the precautionary motive to demand for money¹³. I will be looking at four of these here. The first two give uncertainty virtually no role in determining the rate of interest. The second two profess to have a major role for uncertainty, but it is arguable that they don’t succeed. The aim of the survey is to determine whether there is a plausible explanation of the role of uncertainty that gives it a substantial role in setting the rate of interest. I accept, at least for the sake of the argument, that missing markets and transaction costs might have a role in causing a demand for money, as described in the previous section, and this may be a factor in setting interest rates. Hence the existence of uncertainty is not necessary for interest rates to be positive. The aim here is to see whether there are theoretical reasons for thinking uncertainty makes a substantial difference.

The classic in the literature is Tobin’s 1958 ‘Liquidity Preference as Behaviour Towards Risk’. As the name suggests, in Tobin’s model agents demand money because they are risk averse. Say money has a constant return of 0 and the return on debts is determined by a random variable which is normally distributed with mean $q > 0$ and standard deviation σ . Then agents will spend A on debts, until their risk aversion function b , which is increasing in A and σ , equals q . The theory is that the value an agent places on buying a particular debt is $q - b$, so when $q = b$ there is no point in buying more debts. However, good this is as an analysis of the folk concept of ‘liquidity’, or

¹³ The following survey is indebted to Runde (1994b).

of Keynes's theory of liquidity (and Runde (1994b) argues persuasively that it does badly on both counts) it goes nowhere towards explaining the distinctive role of uncertainty in liquidity preference.

A slightly different slant on liquidity preference is provided by Maclachlan (1993). She defines the liquidity of an asset in terms of the expected transactions costs in trading that asset for some other should the desire arise. This again recalls the need for missing or imperfect markets to explain the importance of money in Keynes's theory. The orthodox explanation for the value of any economic variable is to determine some equations which that value must satisfy in equilibrium, and show that these equations plus some values that we take as exogenous determine the value in question. Maclachlan rightly argues that this approach obliterates the distinction between causation and correlation that a good theory ought preserve. She argues that the only two possible causes for the existence of a positive rate of interest are the desire for liquid assets, in the sense she has defined 'liquid', or the expectation of what Schumpeter called 'super-normal' profits.

While the focus on causation is appropriate, I have doubts about Maclachlan's methodology. She dismisses various potential causes of the rate of interest being positive¹⁴ by showing that they are compatible with the rate being zero. If this were sound it would rule out various causal explanations which we take to be sound. I drop a ball and it falls to the floor. Half way down it is moving at a certain velocity v . Why does it have this velocity? Indeed, why does it have any velocity? The answer to the second question will surely be 'gravitational attraction', and that will be a major part of the answer to the first. Now I throw the ball straight up in the air and watch it climb then fall to earth. At the top of its path its velocity (relative to the earth) is zero. Yet it is under just the same gravitational attraction as it was in the first example. Does this mean that we

¹⁴ For example productivity and time preference.

can't say gravity is a cause of its velocity in the first example? No; gravity is not only a cause, it is the primary cause, and it is a 'positive' cause. The ball does not have the velocity it has despite the best efforts of gravity, as is the case when it is travelling upwards. So on methodological grounds it is no objection to saying F is the cause of interest rates being positive (and taking roughly the particular value they do) to say F might occur and rates might yet be zero.

This methodological diversion is a little irrelevant to our main purpose however. I have accepted that transaction costs might cause a demand for money. The question here is whether uncertainty has any particular causal effects. And on Maclachlan's account the answer would have to be no. She provides no reason for agents in uncertain worlds with imperfect markets to have different attitudes towards liquidity than agents in risky worlds with imperfect markets.

Makowski (1989) defines uncertainty as 'risk plus the possibility of learning'. On his theory uncertainty, of this type, leads to liquidity preference because agents prefer to keep their options open until the new information comes in. Because some markets are missing, in particular markets for trading some investments for others, keeping options open means holding money. Runde (1994b) gives a number of ways we might cash out this notion of uncertainty, some of which are quite similar to our concept of there being multiple reasonable degrees of belief. However, I doubt that uncertainty is really doing any work here, as the following example brings out.

A fair coin is to be tossed and then two fair dice are to be rolled. Say p is true iff the coin falls heads and the dice fall double-six, or if the coin falls tails and the dice do not fall double-six. Assume some agent knows that the coin and dice are fair. Hence she knows that the chance of p is one-half. Since it is true that when an agent knows the chance of A then A is not uncertain¹⁵, p

¹⁵ This follows from my definition of uncertainty and the Principal Principle, but there might be a question as to whether such a conclusion has always been made. Given the quote from Keynes (1937a) defining uncertainty in the previous chapter this seems to be how he is using the term 'uncertainty', and this is

is not uncertain here. Now assume that this agent will learn the result of the coin toss before the market. There is definitely a ‘possibility of learning’ here. Jones and Ostroy (1984) attempt to capture a similar idea to Makowski’s through the variability of beliefs. Again, here the agent’s beliefs about p are very highly variable; she currently believes p to degree $1/2$, but soon she will believe it to nearly degree 1 or 0. And given some simple assumptions about the markets – basically that our agent is barred from selling on bets – she has an incentive to hold money until she learns how the coin falls and then place her bets. So we can get all the crucial features of Makowski’s account without allowing in any uncertainty.

As I stressed earlier, showing this is not enough to show that uncertainty bears no load in Makowski’s story. However, there are reasons for thinking this example is in its essential features perfectly general. What is important in his story is that an agent is (probably) about to learn some information that will be a guide as to what investments are most valuable. Hence she wants to be in the best possible position to take advantage of this information, i.e. hold the most liquid assets. That story seems coherent, and it seems plausible enough to suggest it frequently occurs. However, I didn’t say whether the agent’s initial beliefs were precise or imprecise in that story. Indeed, that seems altogether irrelevant to the story. So whether the initial situation is one of risk or uncertainty doesn’t matter. That is, there is no theoretical reason here for saying uncertainty plays an important role.

standard among theorists who have followed Knight in making a distinction between risk and uncertainty. This point may have been obscured by the fact that many of these theorists held unusual or anachronistic theories of chance (particularly frequentist theories) and that they didn’t refer to chances explicitly, but despite this the assumption is widely shared. In particular, this conclusion seems to be drawn by Hart (1942) and Tintner (1941) – the earliest theorists to suggest representing uncertainty by sets of probability functions – whom both Makowski (1989) and Jones and Ostroy (1984) credit as authorities.

Finally, Paul Davidson (1988, 1991) argues that uncertainty arises whenever agents do not have sufficient knowledge to calculate the numerical probability of an event. He gives this a rather frequentist gloss, but that's not necessary. His idea is that we know what the probability of p is when we know the frequency of p -type events in the past and we know the future will resemble the past in this respect. The latter is cashed out as saying p is governed by an 'ergodic process'. We can replace all this by saying that p is subject to uncertainty whenever we don't know its objective chance, whether or not objective chance ought be analysed by frequentist approaches. Davidson then argues that since for most p we don't have this knowledge¹⁶, we have to adopt 'sensible' approaches like holding money.

Runde (1994b) objects that Davidson's story is incoherent. On Davidson's theoretical story there are only two epistemic states relative to p that are possible. An agent can know the chance of p , in which case their degree of belief is set equal to it, or they are completely uncertain about it. In the latter case there can be no reason for taking some action rather than another. Now the reason that it is 'sensible' to hold money is that we expect money to be liquid. But we don't know the chance of money remaining liquid; whether or not money remains liquid is not determined by an ergodic process. Hence we have no reason for letting that partial belief be a guide to action.

This is a fair criticism, but it can be met by amending the theory rather than by giving it up. On my theory, if an agent knows the chance of p , she will have a precise degree of belief in p . When she doesn't the degree of belief will, in general, be vague but not totally vague. As with Keynes, I have uncertainty come in degrees. This amendment is enough to rescue Davidson. An agent might not know the chance that money will go illiquid in the next short period of time, but she might

¹⁶ His reasoning here is rather dubious. He seems to think that human actions being governed by ergodic processes is inconsistent with the existence of free will. This is a rather strong form of incompatibilism, and even less plausible than the ordinary kind. Nevertheless, we can be confident that agents rarely know the chances of events that are relevant to their wealth.

know enough for it to be reasonable to have a degree of belief which is vague over, say, $[10^{-8}, 10^{-5}]$ in that proposition. And those numbers are small enough that it may still be sensible to hold some money when the expected return on other investments really is vague. So I have to turn to the question of whether it really is sensible to prefer fixed to uncertain returns. Since that is just the question I must answer when looking at whether there is a direct effect of uncertainty that makes people prefer bonds to investments, I turn to that question now.

11.5 Uncertainty and Indecision

As Keynes repeatedly stressed, investment is not like a game of chance where the expected results are known in advance. And this fact is part of the explanation for the extreme instability in investment levels compared to other economic variables¹⁷.

The state of long-term expectation ... does not solely depend on the most probable forecast we can make. It also depends on the *confidence* with which we make this forecast (*GT*: 148).

Businessmen play a mixed game of skill and chance, the average results of which to the players are not known by those who take a hand. If human nature felt no temptation to take a chance, no satisfaction (profit apart) in constructing a factory, a railway, a mine or a farm, there might not be much investment merely as a result of cold calculation (*GT*: 150).

Human decisions affecting the future, whether personal or political or economic, cannot depend on strict mathematical expectation, since the

¹⁷ At *GT*: 103 Keynes quotes Simon Kuznets's research showing that between 1929 and 1932 gross investment levels in the United States fell by over 75% and net investment by over 95%. While nothing as dramatic has subsequently occurred, investment remains one of the most volatile economic indicators.

basis for making such calculations does not exist ... it is our innate urge to activity which makes the wheels go round, our rational selves choosing between the alternatives as best we are able, calculating where we can, but often falling back for our motive on whim or sentiment or chance (*GT*: 162-3).

The liquidity premium ... is partly similar to the risk premium, but partly different; – the difference corresponding to the difference between the best estimates we can make of probabilities and the confidence with which we make them (*GT*: 240).

Our desire to hold money as a store of wealth is a barometer of the degree of our distrust of our own calculations and conventions concerning the future (Keynes 1937a: 116).

As mentioned earlier, the most charitable reading of Keynes is to say he held what I referred to in Chapter 9 as a Horvitz-style decision rule. If the expected return of an investment is vague over $[\alpha, \beta]$ then its 'value' is given by $(1 - \rho)\alpha + \rho\beta$, where $\rho \in [0, 1]$ is a measure of confidence. By the 1937 article he has become more interested in the special case where confidence has collapsed and ρ is approaching 0. This theory would explain all the quotes here, provided we make the safe assumption that 'cold calculation' would only have us spend x on an investment with expected return $[\alpha, \beta]$ when $\alpha \geq x$. In particular any interpretation of the underlying decision theory here will have to give some role to 'whim or sentiment or chance', and I give it a variable, ' ρ '. With this theory I have the extensions needed to avoid Runde's objection to Davidson. I have a continuum of degrees of uncertainty, rather than a raw dichotomy, and I have an explanation of why it is 'sensible' to prefer gambles with known expected returns, at least when ρ is relatively low.

This theory is meant to serve two related purposes. It is meant to show why we might prefer money to debts, even though our best estimate of the expected return of the debts is positive; and, again, it is meant to show why we might prefer debts to investments even when our best estimate of the expected return of the investment is higher. And I think if the decision rule stipulated were plausible, it would show that uncertainty did have an economic effect. In particular, I think it would show both that in times of crises when ρ heads down, the level of investment will decrease even with other things being equal, and that collective action can be justified even when individual action is not¹⁸.

The decision theory doesn't, however, seem plausible. First, there are the technical problems presented for this theory in chapter 9. I won't go over these again, except to note that they gave implausible results whenever we could expect uncertainty to increase in the near future. When this was the case, these theories make recommendations which we know they would reverse whatever information comes in. If it is stipulated that the only possible types of information that will come in will not increase uncertainty, this problem goes away. I leave it to the reader to determine how much such a restriction would limit the applicability of this theory, because I think there is a more serious philosophical objection.

The problem is just the one I raised against Levi's two-step decision theories. What precisely is ρ supposed to represent? If it is some kind of belief, its effects should have been incorporated into the degrees of belief. If it is some kind of desire, its effects should have been incorporated into the evaluation of each of the states. This objection could be avoided, perhaps, if Keynes was trying to argue against the theory that investors just maximise dollar expected returns. It isn't entirely clear who Keynes does think he is arguing against at some points. But if this is his enemy he is fighting a straw man, one who is vulnerable to much simpler objections. Whoever thought that

¹⁸ That is, the government can make sets of investments which are expected to be profitable even though none of the individual investments are expected to be profitable.

all investment is profit driven, that no one ever went into business because they thought it would be fun to run a newspaper? Keynes's only viable opponents here are saying that investors calculate the expected return, in utils, of each possible investment and choose the one whose returns are highest. Now, perhaps, for many investors dollar returns are the most important factor in determining util returns, but they are certainly not the only factor.

If ρ represents something which is neither a belief nor a desire, then it is hard to see what effect it could have on action. Perhaps there are some exceptions to the Humean belief-desire model of explaining agent's actions, for example in actions which are caused by the agent's values, but these exceptions are a million miles away from the kind of cases Keynes considers. And Keynes doesn't seem to appeal to such exemptions. After all, he describes investment decisions made where the 'cold calculations' do not determine what should be done as being made by 'whim or sentiment or chance'. Now whims and sentiments are surely desires, although chance is in a different boat. If he had just said 'chance' here, he may have committed himself to the decision theory I called Caprice; and since there turn out to be good arguments for Caprice, this would have been a prescient move. But he didn't make it; rather Keynes, like Levi, seems to be committed to a kind of mental state which must both be and not be a desire to do the work he requires of it. So even if the technical objections based around Monte Hall type cases are not persuasive because the counterexamples are so fantastic, there are philosophical objections to this position based around its implausible assumptions about mental causation.

11.6 Disquietude

There are some comments from Keynes that suggest this reading is unfair. Rather than having a distinctive decision theory, he perhaps has a distinctive theory about what ought to enter into the decision-theoretic calculations. The standard theory for why there is a demand for insurance is the falling marginal utility of money. Agents purchase insurance, and accept a lower expected dollar return because with insurance their expected util return, at the end of the duration of the

insurance, is higher than if they hadn't purchased. This is the story given in, for example, Friedman and Savage (1952), where the existence of demand for insurance is taken as evidence for the declining marginal utility of money. But there is another reason agents might buy insurance. They might simply feel happier, over the duration of the insured period, knowing that they have insurance and are hence exposed to fewer risks or uncertainties than otherwise. If this is true, then their expected 'wealth' in both dollars and utils at the end of a period might be lower if they insure than if otherwise, but it will be worthwhile because of the benefits during the period. Keynes suggests that this same desire for quietude can cause a demand for money. I presume, though it isn't entirely clear, that this desire should be included within the precautionary motives for holding money.

There are not two separate factors affecting the rate of investment, namely, the schedule of the marginal efficiency of capital [the expected return of investments] and the state of confidence. The state of confidence is relevant because it is one of the major factors determining the former (*GT*: 149).

For the fact that each individual investor flatters himself that his commitment is "liquid" ... calms his nerves and makes him much more willing to run a risk (*GT*: 160).

The possession of actual money lulls our disquietude; and the premium which we require to make us part with money is the measure of the degree of our disquietude (Keynes 1937a: 116).

A liquidity premium ... is not even expected to be rewarded. It is a payment, not for the expectation of increased tangible income at the end

of the period, but for an increase sense of comfort and confidence during the period (Keynes 1938b: 293-4)

This explanation of the demand for certain returns is in some ways conservative and some ways radical. It is conservative because it doesn't immediately change the technical properties of preference. Many of the heterodox theories of preference which have been advocated as responses to the Allais and Ellsberg paradoxes¹⁹ drop such theoretical restrictions as transitivity. A more dramatic response would be to say that the nature of a gamble changes when compared to a certain return. So when choosing between α and a certain million dollars, what matters is not the attractiveness of α , but of α' which is defined as being ' α when one of the choices is a million dollars'. As Sobel (1989: 273-5) points out such theories have to drop the assumption that it is intelligible to make pairwise comparisons between any two goods. We cannot, for example, even compare α' with \$500,000, and hence there is nothing remotely corresponding to the value of α' . By contrast, the theory Keynes appears to be advocating is it least in principle conservative on this front. Agents are still going round maximising expected utility, only now it is expected utility over a period, not at the end of the period.

But it isn't all conservative. Explaining economic decisions in terms of the disquietude of the investor is to discard the distinction between investment and consumption. It was always known that there were some goods that were not comfortably categorised, particularly cars, but this move makes every good in part a consumption good. This isn't important just because some helpful classifications have to be questioned. Rather, its importance flows from its implications for the norms for investment. It is always irrational to make an investment which will incur a sure loss. This principle is used to derive wide-ranging implications for decision-theory. But it is not irrational to make a consumption decision which will result in sure loss at the end of a period in

¹⁹ Many of these are discussed in Chapter 9.

exchange for goods during that period. It is not always irrational to pay ten dollars for a movie ticket, even though this will incur a sure loss in the sense the buyer will surely have less wealth at the end of the movie than if they had not bought the ticket.

Given this, some of the technical complaints I raised against the Horvitz-style decision rule in chapter 9 might miss their target. Indeed, even the philosophical complaint about the role of ρ might be irrelevant. If the expected returns only measure how much various gambles will be worth at the end of the period²⁰, then some desires have not yet been included in our calculations. That is, we can say that ρ represents some desires without being accused of double-counting. So far this all seems to work, and explain the role of uncertainty. Indeed, I think this is the best extension of Keynes's views in this area.

While there seem to be few theoretical objections which can be raised at this point, there is a rather telling empirical objection. The only role given to disquietude in this theory is in deciding between alternatives where the returns on at least one are uncertain. But it seems implausible that disquietude could have this effect, but have no effect when choices are being made between alternatives where at least one is risky. I doubt the feelings of disquiet would be any different were I to have a large fortune riding on a roulette wheel or a baseball game. Disquietude arises because we don't know what will happen; maybe for some people it is greater when we don't know the expected returns, but I doubt it. Again, perhaps there is an explanation for demand for money in the real world to be found here, but uncertainty plays no role in the story, or at best a small cameo.

11.7 Summary

The last two chapters have investigated whether Keynes was correct in thinking that there are interesting economic consequences of his philosophical views. Chapter 10 argued that the best

²⁰ I am assuming the measurements here are conducted in utils.

interpretation of Keynes's theory of probability is the theory advocated in this dissertation. In particular I provide an explanation of *weight* which is consistent with the motivation behind its introduction in the *TP* and use in the *GT*, although not with its explicit formulation in the *TP*. However, this chapter showed that this theory bears no load in economic theorising. However we interpret Keynes's views about the demand for money, or for returns which are more stable in terms of money, any work which is supposed to be done by uncertainty can be done just as well by risk. Indeed there seems to be no situations in which uncertainty is supposed to play a causal role in which risk wouldn't have the same effect. This is not to say that Keynes is wrong here. If anything his theories are strengthened at this point by the fact that he doesn't need a controversial, if correct, theory of uncertainty.

Bibliography

- Arntzenius, Frank and David McCarthy (1997) "The Two-Envelope Paradox and Infinite Expectations" *Analysis* 47: 42-50.
- Arrow, Kenneth (1963) *Social Choice and Individual Values*. New York: Wiley.
- Ayer, Alfred (1936) *Language, Truth and Logic*. London: Gollantz. Second edition 1947.
- References to second edition.
- Ayer, Alfred (1957) The Conception of Probability as a Logical Relation. Reprinted in his *The Concept of a Person*, London: Macmillan Press, 1963, pp 188 - 198.
- Bacchus, Fahiem, Henry Kyburg and Mariam Thalos (1990) "Against Conditionalisation" *Synthese* 85: 475 - 506.
- Barker, Stephen (1995) "Towards a Pragmatic Theory of 'If'" *Philosophical Studies* 79: 185-211.
- Bateman, Brad (1996) *Keynes's Uncertain Revolution*. Ann Arbor: University of Michigan Press.
- Bertrand, Jean (1889) *Calcul des Probabilités*. Paris: Gauthier-Villars. First English edition 1972.
- Bigelow, John and Robert Pargetter (1990) *Science and Necessity*. Cambridge: Cambridge University Press.
- Bigelow, John, John Collins and Robert Pargetter (1993) "The Big Bad Bug: What are the Humean's Chances" *British Journal for the Philosophy of Science* 44: 443-462.
- Black, Max (1967) "Probability" in Paul Edwards *et al* (eds) *The Encyclopaedia of Philosophy*, New York: Macmillan, 8 vols., vol. 6, pp 464-479.
- Blackburn, Simon (1980) "Opinions and Chances" in Mellor (ed), pp 175-196.
- Blackburn, Simon (1984) *Spreading the Word*. Oxford: Clarendon Press.
- Blackburn, Simon (1988) "Attitudes and Contents" *Ethics* 98: 501-517.
- Borel, Emile (1924) "A propos d'un Traité de Probabilités" *Revue Philosophique* 98: 321-336.
- English translation in Kyburg and Smokler (eds) (1964), pp 45-60.

- Broome, John (1995) "The Two-Envelope Paradox" *Analysis* 55: 6-11.
- Carnap, Rudolf (1945) The Two Concepts of Probability. *Philosophy and Phenomenological Research* 5: 513 - 532.
- Carnap, Rudolf (1950) *Logical Foundations of Probability*. Chicago: University of Chicago Press. Second Edition 1962. References to second edition.
- Carnap, Rudolf (1952) *The Continuum of Inductive Methods*. Chicago: University of Chicago Press.
- Carnap, Rudolf (1963) "Replies and Systematic Expositions" in Paul Schilpp (ed) *The Philosophy of Rudolf Carnap*. Illinois: Open Court, pp 859-1016.
- Carnap, Rudolf (1971) Inductive Logic and Rational Decisions. In R. Carnap and R. C. Jeffery (eds) *Studies in Inductive Logic and Probability* Vol. 1, Berkeley: University of California Press, pp 5 - 32.
- Carroll, Lewis (1871) *Through the Looking Glass, and What Alice Found There*. London: Macmillan. Reprinted in *The Annotated Alice*, edited by Martin Gardner, London: Penguin 1960. References to reprint.
- Chisolm, Roderick (1989) "Probability in the Theory of Knowledge" in Marjorie Clay and Keith Lehrer (eds) *Knowledge and Skepticism*. Boulder: Westview.
- Christensen, David (1996) "Dutch-Book Arguments De-Pragmatized" *Journal of Philosophy* 93: 450 - 479.
- Coates, John (1996) *The Claims of Common Sense*. Cambridge: Cambridge University Press.
- Cohen, L. J. (1977) *The Probable and the Provable*. Oxford: Clarendon Press.
- Davidson, Barbara and Robert Pargetter (1985) "In Defence of the Dutch Book Argument" *Canadian Journal of Philosophy* 15: 405 - 424.
- Davidson, Paul (1988) "A Technical Definition of Uncertainty and the Long-Run Non-neutrality of Money" *Cambridge Journal of Economics* 12: 329-338.

- Davidson, Paul (1991) "Is Probability Theory Relevant for Uncertainty? A Post Keynesian Perspective" *Journal of Economic Perspectives* 5: 129-144.
- Davidson, Paul (1994) *Post Keynesian Macroeconomic Theory*. Aldershot: Edward Elgar.
- Davis, John (1994) *Keynes's Philosophical Development*. Cambridge: Cambridge University Press.
- de Finetti, Bruno (1937) "La Prévision; ses lois logiques, ses sources subjectives" *Annales de l'Institut Henri Poincaré* 7: 1-68. First English translation in Kyburg and Smokler (eds) (1964), pp 93-158. References to translation.
- de Finetti, Bruno (1972) *Probability, Induction and Statistics*. New York: Wiley.
- de Finetti, Bruno (1974) *Theory of Probability*. London: Wiley.
- Dempster, Arthur (1967) "Upper and Lower Probabilities induced by a Multi-valued Mapping" *Annals of Mathematical Statistics* 38: 325-339.
- Dempster, Arthur (1968) "A Generalisation of Bayesian Inference" *Journal of the Royal Statistical Society Series B* 30: 205-247.
- Dempster, Arthur (1988) "Probability, Evidence and Judgment" (with discussion) in D. Bell, H. Raiffa and A. Tversky (eds) *Decision Making*. Cambridge: Cambridge University Press. pp 283-298.
- Dummett, Michael (1959) "Truth" *Proceedings of the Aristotelian Society* 59: 141-62.
Reprinted in his *Truth and Other Enigmas*, London: Duckworth, 1978, pp 1 - 19.
References to reprint.
- Dummett, Michael (1975) "Wang's Paradox" *Synthese* 30: 301-324.
- Dummett, Michael (1977) *Elements of Intuitionism*. Oxford: Clarendon Press.
- Edwards, W. H. Lindman and L. J. Savage (1963) "Bayesian Statistical Inference for Psychological Research" *Psychological Review* 70: 193-242.

- Fagin, R. and J. Halpern (1991). A New Approach to Updating Beliefs. In P. Bonissone, M. Henrion, L. Kanal and J. Lemmer (eds) *Uncertainty in Artificial Intelligence 6*. Amsterdam: North-Holland, pp 347-374.
- Fine, Kit (1975) "Vagueness, Truth and Logic" *Synthese* 30: 265-300.
- Fine, Terrence (1973) *Theories of Probability: An Examination of Foundations*. New York: Academic Press.
- Fine, Terrence (1983) "Foundations of Probability" in S. Kotz and N. L. Johnson (eds.) *The Encyclopaedia of Statistical Sciences* Vol. 3, New York: Wiley.
- Fodor, Jerry and Ernest Lepore (1996) "What Cannot be Valuated Cannot be Valuated, and it Cannot be Supervaluated Either" *Journal of Philosophy* 93: 516-535.
- Friedman, Milton and Leonard Savage (1952) "The Expected Utility Hypothesis and the Measurability of Utility" *Journal of Political Economy* 60: 463-474.
- Gardner, Martin (1961) *The Second Scientific American Book of Mathematical Puzzles and Diversions*. New York: Simon & Schuster.
- Geach, Peter (1965) "Assertion" *Philosophical Review* 74: 449-465.
- Gibbard, Alan (1990) *Wise Choices, Apt Feelings*. Oxford: Clarendon Press.
- Gilboa, Itzhak and David Schmeidler (1993) "Updating Ambiguous Beliefs" *Journal of Economic Theory* 59: 33-49.
- Gillies, Donald (1988) "Keynes as a Methodologist" *British Journal for the Philosophy of Science* 39: 117-29.
- Gillies, Donald (1991) "Intersubjective Probability and Confirmation Theory" *British Journal for Philosophy of Science* 42: 513- 533.
- Gillman, Leonard (1992) "The Car and the Goats" *American Mathematical Monthly* 99: 3-8.
- Gödel, Kurt (1931) "Über Formal Unentscheidbare Sätze der *Principia Mathematica* und Verwandter Systeme, I" *Monatshefte für Mathematik und Physik* 38: 173-198.
- Good, I. J. (1950) *Probability and the Weighing of Evidence*. London: Griffin.

- Good, I. J. (1952) "Rational Decisions" *Journal of the Royal Statistical Society Series B* 14, 107 - 114. Reprinted in his *Good Thinking*, Minneapolis: University of Minnesota Press, 1983, pp 3 - 14.
- Good, I. J. (1962) "Subjective Probability as the Measure of a Non-Measurable Set" in E. Nagel, P. Suppes and A. Tarski (eds.) *Logic, Methodology and Philosophy of Science*. Stanford: Stanford University Press, pp 319-329.
- Goodman, Nelson (1947) "The Problem of Counterfactual Conditionals" *Journal of Philosophy* 44: 113-128.
- Goodman, Nelson (1954) *Fact, Fiction and Forecast*. Cambridge, MA.: Harvard University Press.
- Green, Mitchell and Christopher Hitchcock (1994) "Reflections on Reflection: van Fraassen on Belief" *Synthese* 98: 297-324.
- Hacking, Ian (1975) *The Emergence of Probability*. Cambridge: Cambridge University Press.
- Hájek, Alan (1998?) "Objecting Vaguely to Pascal's Wager" To be published?
- Hale, Bob (1993) "Can There be a Logic of Attitudes" in John Haldane and Crispin Wright (eds) *Reality, Representation and Projection*. New York: Oxford University Press, pp 337-364.
- Hall, Ned (1994) Correcting the Guide to Objective Chance. *Mind* 103: 504 - 517.
- Hammond, Peter (1988) "Consequentialist Foundations for Expected Utility Theory" *Theory and Decision* 25: 25-78.
- Hansson, Bengt (1988) "Risk Aversion as a Problem of Conjoint Measurement" in P. Gärdenfors and N. E. Sahlin (eds.) *Decision, Probability and Utility*. Cambridge: Cambridge University Press, pp 117-136.
- Harcourt, Geoff and P. A. Riach (1997) *A 'Second Edition' of the General Theory*. 2 Vols. London: Routledge.
- Harman, Gilbert (1983) "Problems with Probabilistic Semantics" In A. Orenstein and R. Stern (eds) *Developments in Semantics*, New York: Haven, pp 243 - 247.

-
- Hart, A. G. (1942) "Risk, Uncertainty and the Unprofitability of Compounding Probabilities" in O. Lange, F. McIntyre and T. O. Yntema (eds) *Studies in Mathematical Economics and Econometrics*. Chicago: University of Chicago Press.
- Hausman, Daniel (1991) *The Inexact and Separate Science of Economics*. Cambridge: Cambridge University Press.
- Hellman, Geoffrey (1997) "Bayes and Beyond" *Philosophy of Science* 64: 191-221.
- Hempel, Carl (1965) *Aspects of Scientific Explanation*. New York: Free Press.
- Hicks, John (1967) *Critical Essays in Monetary Theory*. Oxford: Clarendon Press.
- Horgan, Terence (1995) "Let's Make a Deal" *Philosophical Papers* 24: 209-225.
- Horwich, Paul (1994) "The Essence of Expressivism" *Analysis* 54: 19-20.
- Howson, Colin (1993) "Dutch Book Arguments and Consistency" *Proceedings of the Philosophy of Science Association* 2: 161-168.
- Howson, Colin and Peter Urbach (1989) *Scientific Reasoning*. La Salle: Open Court.
- Jackson, Frank and Robert Pargetter (1976) "A Modified Dutch Book Argument" *Philosophical Studies* 29, 403 - 407.
- Jaffray, J. (1994) Decision Making with Belief Functions. In Yager, Fedrizzi and Kacprzyk (eds) pp. 331-352.
- Jeffrey, Richard (1983a) "Bayesianism with a Human Face", in J. Earman (ed.) *Testing Scientific Theories*, Minneapolis: University of Minnesota Press.
- Jeffrey, Richard (1983b) *The Logic of Decision*. Chicago: University of Chicago Press, second edition. First edition New York: McGraw Hill, 1965.
- Jeffrey, Richard (1987) "Indefinite Probabilities: Reply to Levi" *Philosophy of Science* 54: 587-591.
- Jeffrey, Richard (1991) "Radical Probabilism" in Enrique Villanueva (ed) *Rationality in Epistemology*. Atascadero: Ridgeview.

- Jeffrey, Richard (1992) *Probability and the Art of Judgement*. Cambridge: Cambridge University Press.
- Jones, R. and J. M. Ostroy (1984) "Flexibility and Uncertainty" *Review of Economic Studies* 51: 13-32.
- Kahnemann, Daniel and Amos Tversky (1979) "Prospect Theory: An Analysis of Decision Under Risk" *Econometrica* 47: 263 - 291.
- Kaplan, Mark (1996) *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.
- Keynes, John Maynard (1910) "Great Britain's Foreign Investments" *New Quarterly* February (1910); *C.W.* XV 46-7
- Keynes, John Maynard (1921a) *Treatise on Probability*. London: Macmillan; *C.W.* VIII.
- Keynes, John Maynard (1921b) "The State and Finance"; *C. W.* XVI 295-307.
- Keynes, John Maynard (1931) Review of *Foundations of Mathematics* by Frank Ramsey, *The New Statesman and Nation* 2: 407; *C.W.* X: 336-9.
- Keynes, John Maynard (1936) *The General Theory of Employment, Interest and Money*. London: Macmillan; *C.W.* VII.
- Keynes, John Maynard (1937a) "The General Theory of Employment", *Quarterly Journal of Economics* 51: 209-23; *C.W.* XIV 109-23.
- Keynes, John Maynard (1937b) "The 'Ex Ante' Theory of the Rate of Interest" *Economic Journal* 47: 663-668. *C. W.* XIV 215-223.
- Keynes, John Maynard (1938a) Letter to Hugh Townshend dated 27 July. *C.W.* XIV 288-9.
- Keynes, John Maynard (1938b) Letter to Hugh Townshend dated 7 December. *C.W.* XIV 293-4.
- Keynes, John Maynard (1938c) "My Early Beliefs"; *C. W.* X: 433-450.
- Keynes, John Maynard (1971-89) *The Collected Writings of John Maynard Keynes* [*C.W.*] ed. D.E. Moggridge, Vols I - XXX, London: Macmillan. (All references to the *C.W.* reprints).
- Knight, Frank (1921) *Risk, Uncertainty and Profit*. Chicago: University of Chicago Press.

- Kolmogorov, A. N. (1933) *Foundations of the Theory of Probability*. First English edition New York: Chelsea Publishing Company, 1950. References to 1950 edition.
- Koopman, Bernard (1940) "The Bases of Probability" *Bulletin of the American Mathematical Society* 46: 763-774.
- Kripke, Saul (1965) "Semantical Analysis of Intuitionistic Logic, I" in M. Dummett and J. Crossley (eds.) *Formal Systems and Recursive Functions*. Amsterdam: North-Holland.
- Kripke, Saul (1972) "Naming and Necessity" In Donald Davidson and Gilbert Harman (eds) *Semantics of Natural Language*. Dordrecht: Reidel.
- Kripke, Saul (1975) "Outline of a Theory of Truth" *Journal of Philosophy* 72: 690-716.
- Kripke, Saul (1982) *Wittgenstein on Rules and Private Language*, Oxford: Blackwell.
- Kvanvig, J. (1994) "A Critique of Van Fraassen's Voluntaristic Epistemology" *Synthese* 98: 325 - 348.
- Kyburg, Henry (1961) *Probability and the Logic of Rational Belief*. Middletown: Wesleyan University Press.
- Kyburg, Henry (1968) "Bets and Beliefs" *American Philosophical Quarterly* 5: 54-63.
Reprinted in his (1983b) pp 63-78. References to reprint.
- Kyburg, Henry (1974) *The Logical Foundations of Statistical Inference*. Dordrecht: Reidel.
- Kyburg, Henry (1978) "Subjective Probability: Criticisms, Reflections and Problems" *Journal of Philosophical Logic* 7: 157-180. Reprinted in his (1983b), pp 79-98. References to reprint.
- Kyburg, Henry (1983a) "The Reference Class" *Philosophy of Science* 50: 374-397.
- Kyburg, Henry (1983b) *Epistemology and Inference*. Minneapolis: University of Minnesota Press.
- Kyburg, Henry (1990) *Science and Reason* New York: Oxford.
- Kyburg, Henry and E. Smokler (eds) (1964) *Studies in Subjective Probability*. Huntington: Krieger.

-
- Lawson, Tony (1985) "Uncertainty and Economic Analysis" *Economic Journal* 95: 909-927.
- Levi, Isaac (1974) "On Indeterminate Probabilities" *Journal of Philosophy* 71: 391-418.
- Levi, Isaac (1980) *The Enterprise of Knowledge*. Cambridge, MA.: MIT Press.
- Levi, Isaac (1986) *Hard Choices*. Cambridge: Cambridge University Press.
- Levi, Isaac. (1987) "The Demons of Decision" *Monist* 70: 193 - 211.
- Lewis, David (1976a) "The Paradoxes of Time Travel" *American Philosophical Quarterly* 13: 145-152.
- Lewis, David (1976b) "Probability of Conditionals and Conditional Probability" *Philosophical Review* 85: 297-315.
- Lewis, David (1979a) "Scorekeeping in a Language Game" *Journal of Philosophical Logic* 8: 339 - 359. Reprinted in his (1983), pp 233 - 249.
- Lewis, David (1979b) "Attitudes *De Dicto* and *De Se*" *Philosophical Review* 88: 513 - 543.
- Lewis, David (1980) "A Subjectivist's Guide to Objective Chance" In R. C. Jeffrey (ed) *Studies in Inductive Logic and Probability* Vol. 2, Berkeley: University of California Press, Reprinted with additional postscripts in his (1986b), pp 83 - 132.
- Lewis, David (1981) "Causal Decision Theory" *Australasian Journal of Philosophy* 59: 5-30.
- Lewis, David (1986a) "Probabilities of Conditionals and Conditional Probabilities" *Philosophical Review* 95: 581-589.
- Lewis, David (1986b) *Philosophical Papers* Vol. 2, New York: Oxford University Press.
- Lewis, David (1988) "Desire as Belief" *Mind* 97: 323-332.
- Lewis, David (1994) "Humean Supervenience Debugged" *Mind* 103: 473 - 490.
- Lewis, David (1996) "Elusive Knowledge" *Australasian Journal of Philosophy* 74: 549 -567.
- Lewis, David (1997) "Finkish Dispositions" *Philosophical Quarterly* 47, 143 - 158.
- Lowe, E J (1996) "Conditional Probability and Conditional Beliefs" *Mind*: 105, 603 - 615.
- Maclachlan, Fiona (1993) *Keynes's General Theory of Interest: A Reconsideration*. London: Routledge.

-
- Maher, Patrick (1992) "Diachronic Rationality" *Philosophy of Science* 59: 120 - 141.
- Maher, Patrick (1993) *Betting on Theories*. Cambridge: Cambridge University Press.
- Maher, Patrick (1997) "Depragmatised Dutch Book Arguments" *Philosophy of Science* 64: 291-305.
- Makowski, Louis (1989) "Keynes's Liquidity Preference Theory: A Suggested Reinterpretation" in Frank Hahn (ed) *The Economics of Missing Markets, Information and Games*. Oxford: Clarendon Press.
- Martin, C. B. (1994) "Dispositions and Conditionals" *Philosophical Quarterly* 44: 1 - 8.
- McCauley, James (1993) *Everything Linguists Have Ever Wanted to Know About Logic (But Were Afraid to Ask)* Chicago: University of Chicago Press. Second edition. First edition 1981.
- McDermott, Michael (1996) "On the Truth-Conditions of Certain 'If'-Sentences" *Philosophical Review* 105: 1-37.
- Mellor, Hugh (ed) (1980) *Prospects for Pragmatism*. Cambridge: Cambridge University Press.
- Morgan, Charles and Edward Mares (1995) "Conditionals, Probability and Non-Triviality" *Journal of Philosophical Logic* 24: 455-467.
- Morgan, Charles and Hughes Leblanc (1983a) "Probabilistic Semantics for Formal Logic" *Notre Dame Journal of Formal Logic* 24: 161-180.
- Morgan, Charles and Hughes Leblanc (1983b) "Probability Theory, Intuitionism, Semantics and the Dutch Book Argument" *Notre Dame Journal of Formal Logic* 24: 289-304.
- Moser, Paul and D. Hudson Mulder (1994) "Probability in Rational Decision-Making" *Philosophical Papers* 23: 109-128.
- Nolan, Daniel (1997) "Quantitative Parsimony" *British Journal for the Philosophy of Science* 48: 329-342.
- Norton, John (1998) "When the Sum of Our Expectations Fails Us" *Pacific Philosophical Quarterly* 79: 34-58.

- O'Donnell, Rod (1989) *Keynes: Philosophy, Economics and Politics*. London: Macmillan Press.
- O'Donnell, Rod (1991) "Keynes on Probability, Expectations and Uncertainty" in O'Donnell (ed) (1991) *Keynes as Philosopher - Economist* London: Macmillan, pp 3-60.
- Paris, J. B. (1994) *The Uncertain Reasoner's Companion*. Cambridge: Cambridge University Press.
- Pearl, Judea (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo: Morgan Kaufman.
- Ponsonnet, Jean-Marc (1996) "The Best and the Worst in G. L. S. Shackle's Decision Theory" in Schmidt (ed) 169-196.
- Popper, Karl (1933) *Logic of Scientific Discovery*. London: Routledge. First English edition 1959.
- Popper, Karl (1959) "The Propensity Interpretation of Probability" *British Journal for the Philosophy of Science* 10: 25-42.
- Price, Huw (1984) "Mellor, Chance and the Single Case" *British Journal for the Philosophy of Science* 35: 11-23.
- Raiffa, H. (1968) *Decision Analysis: Introductory Lectures on Choice Under Uncertainty*. Reading, MA.: Addison-Wesley.
- Ramsey, Frank (1926a) "The Foundations of Mathematics" in his (1931), pp 62-81.
- Ramsey, Frank (1926b) "Truth and Probability" in his (1931), pp 156-198.
- Ramsey, Frank (1928) "Chance" in his (1931), pp 206 - 211.
- Ramsey, Frank (1929a) "Theories" in his (1931), pp 212-236.
- Ramsey, Frank (1929b) "Probability and Partial Belief" in his (1931), pp 256-7.
- Ramsey, Frank (1931) *The Foundations of Mathematics*. Edited by R. B. Braithwaite, London: Routledge.
- Rogers, Colin and T. K. Rymes (1997) "Keynes's Monetary Theory of Value and Modern Banking" in Harcourt and Riach (eds), vol. 1, pp 304-323.

-
- Runde, Jochen (1990) "Keynesian Uncertainty and the Weight of Arguments" *Economics and Philosophy* 6: 275-93.
- Runde, Jochen (1994a) "Keynes After Ramsey: In Defence of 'A Treatise on Probability'" *Studies in the History and Philosophy of Science* 25: 97-124.
- Runde, Jochen (1994b) "Keynesian Uncertainty and Liquidity Preference" *Cambridge Journal of Economics* 18: 129-144.
- Russell, Bertrand (1905) "On Denoting" *Mind* 14: 479-493.
- Russell, Bertrand (1922) "Review of John Maynard Keynes, *A Treatise on Probability*" *Mathematical Gazette* 11: 119-125.
- Russell, Bertrand (1923) "Vagueness" *Australasian Journal for Psychology and Philosophy* 1, 84-92.
- Russell, Bertrand (1940) *An Inquiry Into Meaning and Truth*. London: Allen and Unwin.
- Russell, Bertrand (1948) *Human Knowledge: Its Scope and Limits*. London: Allen and Unwin.
- Sainsbury, Mark (1988) *Paradoxes*. Cambridge: Cambridge University Press.
- Saul, Jennifer (1997) "Substitution and Simple Sentences" *Analysis*, 57: 102 - 109.
- Savage, Leonard (1954) *The Foundations of Statistics*. New York: John Wiley.
- Savage, Leonard (1964) "The Foundations of Statistics Reconsidered" In Kyburg and Smokler (eds) (1964), pp 173-188.
- Savage, Leonard (1967a) "Difficulties in the Theory of Personal Probability" *Philosophy of Science* 34: 305-310.
- Savage, Leonard (1967b) "Implications of Personal Probability for Induction" *Journal of Philosophy* 64: 593-607.
- Schick, Frederick (1986) "Dutch Bookies and Money Pumps" *Journal of Philosophy* 83, 112 - 119.
- Schmeidler, David (1989) "Subjective Probability and Expected Utility Without Additivity" *Econometrica* 57: 571-589.

-
- Schmidt, Christian (ed) (1996) *Uncertainty in Economic Thought*, Cheltham, Edward Elgar.
- Scott, Alexander and Michael Scott (1997) "What's In the Two-Envelope Paradox" *Analysis* 57: 34-41.
- Seidenfeld, Teddy (1984) Review of Paul Horwich's *Probability and Evidence*. *Philosophical Review* 93 474-483.
- Seidenfeld, Teddy (1994) "When Normal and Extensive Form Decisions Differ" in *Logic, Methodology and Philosophy of Science*, edited by Dag Prawitz, Brian Skyrms and Dag Westerståhl, Amsterdam: Elsevier, pp 451-463.
- Shackle, George (1949) *Expectation in Economics*. Cambridge: Cambridge University Press.
- Shackle, George (1961) *Decision, Order and Time in Human Affairs*. Cambridge: Cambridge University Press.
- Shackle, George (1972) *Epistemics and Economics: A Critique of Economic Doctrines*. Cambridge: Cambridge University Press.
- Shafer, Glenn (1976) *A Mathematical Theory of Evidence*. Princeton: Princeton University Press.
- Shafer, Glenn (1981) "Constructive Probability" *Synthese* 48: 1 - 60.
- Shapiro, Nina (1997) "Imperfect Competition and Keynes" in Harcourt and Riach (eds), vol. 1, pp 83-92.
- Skidelsky, Robert (1983) *John Maynard Keynes*. Vol. I: *Hopes Betrayed, 1883-1920*. London: Macmillan.
- Skidelsky, Robert (1992) *John Maynard Keynes*. Vol. II: *The Economist as Saviour, 1920-1937*. London: Macmillan.
- Slutsky, E. (1915) "On the Theory of the Budget of the Consumer" *Giornale delgi Economisti* 51: 1 - 26. Translated in K. Boulding and G. Stigler (eds.) *Readings in Price Theory* Chicago: Irwin, pp. 27-56.

- Smets, Phillippe (1994) "What is Dempster-Shafer's Theory of Evidence?" In Yager, Fedrizzi and Kacprzyk (eds) pp 5-34.
- Smith, Cedric A. B. (1961) "Consistency in Statistical Inference and Decision" *Journal of the Royal Statistical Society* Series B 23: 1 - 37.
- Sobel, J. Howard (1989) "Utility Theory and the Bayesian Paradigm" *Theory and Decision* 26: 263-293.
- Strat, T. (1990) "Decision analysis using Belief Functions" *International Journal of Approximative Reasoning* 4: 391-417.
- Strawson, Peter (1950) "On Referring" *Mind* 59: 320-344.
- Strevens, Michael (1995) "A Closer Look at the 'New' Principle" *British Journal for the Philosophy of Science* 46: 545 - 561.
- Suppes, Patrick (1974) "The Measurement of Belief" *Journal of the Royal Statistical Society* Series B 36: 160-191.
- Teller, Paul (1973) "Conditionalization and Observation" *Synthese* 26: 218 - 258.
- Thau, Michael (1994) "Undermining and Admissibility" *Mind* 103: 491 - 503.
- Tintner, Gerhard (1941) "The Theory of Choice Under Subjective Risk and Uncertainty" *Econometrica* 9: 298-304.
- Tobin, James (1958) "Liquidity Preference as Behaviour Towards Risk" *Review of Economic Studies* 25: 65-86.
- Tooley, Michael (1987) *Causation: A Realist Approach*. Oxford: Oxford University Press.
- van Fraassen, Bas (1966) "Singular Terms, Truth-Value Gaps and Free Logic" *Journal of Philosophy* 66: 481-95.
- van Fraassen, Bas (1984) "Belief and the Will" *Journal of Philosophy* 81: 235 - 256.
- van Fraassen, Bas (1989) *Laws and Symmetry*. Oxford: Clarendon Press.
- van Fraassen, Bas (1990) "Figures in a Probability Landscape" in J. Dunn and A. Gupta (eds) *Truth or Consequences*. Amsterdam: Kluwer, pp 345-356.

-
- van Fraassen, Bas (1995) "Belief and the Problem of Ulysses and the Sirens" *Philosophical Studies* 77: 7 - 37.
- von Kries, J. (1886) *Die Principien der Wahrscheinlichkeitsrechnung. Eine logische Untersuchung*. Frieberg.
- Walley, Peter (1991) *Statistical Reasoning with Imprecise Probabilities* London: Chapman & Hall.
- Wang, Hao (1987) *Reflections on Gödel*. Cambridge, MA: MIT Press.
- Williams, Peter (1976) Indeterminate Probabilities. In M. Przelecki, K. Szaniawski and R. Wojcicki (eds) *Formal Methods in the Methodology of the Empirical Sciences*. Dordrecht: Reidel, pp 229-246.
- Williams, Peter (1978) "On a New Theory of Epistemic Probability" *British Journal for the Philosophy of Science* 29: 375-387.
- Williamson, Timothy (1994) *Vagueness*. London: Routledge.
- Wittgenstein, Ludwig (1953) *Philosophical Investigations*. London: Routledge.
- Wolfenson, Marco and Terrence Fine (1982) "Bayes-Like Decision Making with Upper and Lower Probabilities" *Journal of the American Statistical Society* 72: 80-88.
- Yager, R. , M. Fedrizzi and J. Kacprzyk (eds) (1994) *Advances in the Dempster-Shafer Theory of Evidence*. New York, NY.: John Wiley.
- Zabell, S. L. (1991) "Ramsey, Truth and Probability" *Theoria* 57: 211-238.
- Zadeh, Lofti (1978) "Fuzzy Sets as a Basis for a Theory of Probability" *Fuzzy Sets and Systems* 1: 3-28.

Name Index

—A—

Allais, M., 130, 341
 Arntzenius, F., 116, 118
 Arrow, K., 280, 291, 292, 293
 Ayer, A. J., 2, 13, 14, 15, 16, 23–32, 147, 312,
 313

—B—

Bacchus, F., 51
 Barker, S., 30
 Bateman, B., 14, 203, 298, 311, 312, 313, 314
 Bertrand, J., 83, 84, 195
 Bigelow, J., ix, 146, 165
 Black, M., 22
 Blackburn, S., 23, 27, 28, 31
 Borel, E., 72, 198
 Broome, J., 116

—C—

Carnap, R., vi, 1, 2, 16–26, 33, 34, 134, 135,
 136, 139, 140, 142, 145, 156, 160, 166, 204,
 210, 213, 214, 215, 303, 306
 Chase, J., ix
 Chisholm, R., 147
 Christensen, D., 40, 41, 127, 128
 Clinton, W., 237, 238, 239, 241
 Cohen, L. J., 242
 Collins, J., 165

—D—

Davidson, B., 53, 54
 Davidson, P., 319, 320, 334, 335, 337
 Davis, J., 14, 311
 de Finetti, B., 6, 22, 23, 24, 32, 33, 34, 38, 50,
 112, 113
 Dempster, A., 85, 87, 101, 102, 109, 242, 291,
 292, 295
 Dodgson, C., 266

Dummett, M., 40, 179, 183, 184, 185, 245, 266

—E—

Ellsburg, D., 130, 341

—F—

Fagin, R., 92, 93, 94, 105, 106, 107
 Fine, K., 168, 174, 183, 187
 Fine, T., 16, 98, 299
 Fodor, J., 185–93
 Friedman, M., 339

—G—

Gardner, M., 83, 84
 Geach, P., 29
 Gibbard, A., 27, 272
 Gilboa, I., 108, 109, 276
 Gillies, D., 14, 145, 311
 Gillman, L., 119
 Glymour, C., 231
 Gödel, K., 310
 Good, I. J., 39, 56, 78, 79, 80
 Goodman, N., 19, 154
 Grice, H. P., 143

—H—

Hacking, I., 146
 Hájek, A., 5, 276
 Hale, B., 31
 Halpern, J., 92, 93, 94, 105, 106, 107
 Hansson, B., 91
 Harman, G., 40, 47, 244
 Harper, W., 272
 Hart, A. G., 79, 80, 333
 Hausman, D., 280
 Hempel, C., 233, 315
 Hicks, J., 325
 Holton, R., ix
 Horgan, T., 296

Horwich, P., 29
 Howson, C., 19, 23, 24, 68, 127, 128, 204, 214,
 215, 216, 217
 Humberstone, I. L., ix

—J—

Jaffray, J., 278, 279, 323
 Jeffrey, R., 19, 21, 78, 93, 94, 196, 225, 226,
 230, 231, 232, 292, 298
 Jones, R., 334

—K—

Kahnemann, D., 38
 Kaplan, M., 127, 128, 157, 158, 213
 Keynes, J. M., vi, vii, 1, 2, 14, 16, 19, 23, 26,
 34, 35, 78, 134, 153, 156, 161, 196–211, 218,
 270, 297–343
 Knight, F., 78, 79, 80, 307, 333
 Kolmogorov, A. N., 4, 5, 246
 Koopman, B., 56, 71
 Kripke, S., 250, 251, 252, 253, 254, 255, 257,
 258, 262, 264, 326
 Kuznets, S., 336
 Kvanvig, J., 40
 Kyburg, H., 8, 10, 20, 23, 26, 33, 51, 232–41,
 304

—L—

Lawson, T., 307
 Leblanc, H., 246, 247, 248
 Lepore, E., 185–93
 Levi, I., 49, 50, 78, 218–26, 279, 280, 281, 282,
 283, 284, 285, 286, 287, 291, 292, 295, 322,
 323, 338, 339
 Lewis, D., 5, 26, 28, 41, 53, 84, 135, 136, 137,
 138, 141, 149, 150, 151, 156–65, 180, 215,
 247, 257, 265, 272
 Locke, J., 147
 Lowe, E. J., 141

—M—

Maclachlan, F., 331, 332
 Maher, P., 49, 50, 127–33, 157, 158

Makowski, L., 333, 334
 Mares, E., 246, 247, 248
 Martin, C. B., 53
 McCarthy, D., 116, 118
 McCauley, J., 63
 McDermott, M., 185
 Mellor, D. H., 36
 Menzies, P., 164
 Morgan, C., 246, 247, 248
 Moser, P., 295
 Mulder, D. H., 295

—N—

Nolan, D., ix, 144
 Norton, J., 275

—O—

O'Donnell, R., 156, 307, 308, 309
 Oppy, G., ix
 Ostroy, J. M., 334

—P—

Pargetter, R., 53, 54, 146, 165
 Paris, J. B., 106
 Pearl, J., 86
 Pettit, P., ix
 Ponsonnet, J-M., 245
 Popper, K., 2, 11, 12, 16, 153
 Price, H., 28, 29, 30, 172

—Q—

Quine, W. V. O., 141

—R—

Ramsey, F. P., 2, 21, 28, 35–41, 49, 55, 56, 147,
 176, 202–13, 265, 266, 270, 284, 297, 298,
 306, 309, 310
 Rogers, C., 319
 Runde, J., 14, 203, 298, 299, 305, 330, 331,
 333, 335, 337
 Russell, B., 8, 9, 19, 20, 91, 145, 181, 182, 303,
 309, 310
 Rymes, T. K., 319

—S—

Sainsbury, M., 173, 178
 Saul, J., 143
 Savage, L., 22, 24, 25, 38, 50, 56, 127, 339
 Schick, F., 48, 49
 Schmeidler, D., 89, 90, 91, 92, 108, 109, 242, 276
 Schmidt, C., 78
 Scott, A., 117
 Scott, M., 117
 Seidenfeld, T., 132, 274, 286
 Shackle, G., 242, 245, 265
 Shafer, G., 52, 58, 59, 80, 81–88, 101, 102, 106, 109, 111, 127, 128, 145, 146, 242, 243, 244, 246, 291, 295
 Shapiro, N., 318
 Skidelsky, R., 298, 305, 310
 Skyrms, B., 272
 Slutsky, E., 43
 Smets, P., 291
 Smith, A., 43
 Smith, C. A. B., 74, 76, 78, 114
 Sobel, J. H., 272, 341
 Strat, T., 278, 279, 323
 Strawson, P., 184, 185
 Suppes, P., 98, 100

—T—

Talmont-Kaminski, K., ix
 Tarski, A., 168

Tarski, A., 167, 173, 191
 Thalos, M., 51
 Tintner, G., 79, 333
 Tobin, J., 331
 Tooley, M., 20
 Townshend, H., 307
 Tversky, A., 38

—U—

Urbach, P., 19, 23, 24, 127, 128, 204, 214, 215, 216, 217

—V—

van Fraasen, B., 46
 van Fraassen, B., 9, 18, 22, 41, 46, 47, 49, 68, 69, 74, 78, 107, 108, 191, 224, 226–30, 291
 von Kries, J., 195

—W—

Walley, P., 59, 74, 98, 99, 100, 110–21
 Whitehead, A., 91, 303
 Williams, P., 74, 76, 80, 114, 174, 286
 Williamson, T., 166, 173, 174, 176, 178, 182, 183, 184, 190, 198
 Wittgenstein, L., 209
 Wolfenson, M., 98, 100

—Z—

Zabell, S. L., 203
 Zadeh, L., 242

Subject Index

- 3—
- 3-coherence, 98
- A—
- Absolutism, 154
- Addition, 127–33, 243
- Additivity, 5, 112–13, 150, 223
- countable, 5, 6, 114
 - finite, 5
- Ambiguity theory, 140–53
- Analyses of probability
- conventionalist, 13, 311–17
 - frequentist, 6
 - intersubjective. *See* conventionalist
 - logical, 16–21
 - modal frequency, 10
 - necessitarian, 138, 202
 - partial entailment. *See* logical
 - propensity, 11–12
 - reasonable degree of belief, 1, 134–65
 - subjectivist, 21–34
 - syntactic. *See* logical
- Anti-realism, 129, 245
- Arguments for imprecision
- ignorance, 72, 195–96
 - introspection, 71
 - Levi's, 218–23
 - rational disagreement, 73, 196–97
 - vagueness, 197–99
 - van Fraassen's, 226
- Arrow's Theorem, 280, 292
- B—
- Basic Chance Principle, 165
- BCP. *See* Basic Chance Principle
- Belief
- is degree of belief one, 62, 156–59
- Belief functions. *See* Shafer functions
- Bels.* *See* Shafer functions
- Betting analysis, 35, 37–41, 52–54, 72, 76, 98, 100, 127, 133, 157
- and constructivism, 244–45
 - and incoherent agents, 41
 - restrictions on, 44
- C—
- c*-functions, 17
- Chance, 136–38, 140–53, 159–65
- non-numerical, 149, 319
- Coherence, 62
- closure, 62, 65
 - consistency, 62, 65
 - dynamic, 65
 - justification, 213
 - Kyburg's objection, 234
- Complete monotonicity. *See* monotonicity
- Conditionalisation, 33, 58, 59, 65, 68, 69, 71, 77, 104, 105, 106, 107, 113, 149, 216, 219, 223, 224, 227–30, 276, 294, 295
- Conglomerability, 100, 110–21, 275, 277
- defined, 112
 - finite, 103
 - restricted, 275
- Constructivist Probability, 130, 242–69
- and anti-realism, 245
 - and betting analysis, 244
 - and Equivalence Analysis, 246
 - axioms for, 249
 - falsum, 248
 - ignorance, 243
 - intuitionist probability, 256–60
 - Kripke trees, 250–60
 - objections, 264
 - updating, 260
 - verificationism, 243
- Contextualism, 156–59
- Continuity, 225, 272
- Conventionalism, 154–56

Conversational score, 141

Convexity, 94, 225–26

—D—

Debts, 321

Decision rules

Caprice, 288–92

conservatism, 286–88

global dominance, 276

Horvitz, 278, 337

Levi's, 279–86, 338

Maxi, 277–79

maximin, 276–77, 292

Decision theory

defined, 270

Degrees of belief, 55–165

and behaviour, 37

imprecise, 71–121

introspective, 36

precise, 2, 58

qualitative and quantitative, 56

rational-valued, 58

Dempster-Shafer functions. *See* Shafer

functions

Dilation, 274

Dispositions, 15, 27, 36, 39, 51, 53, 54, 112,

133, 158, 198, 200, 227, 231, 296

finkish, 53, 158

Disquietude, 339–42

Dutch Book Argument, 2, 22, 35, 40, 42–52, 57,

60, 65, 68, 69, 75, 127–33, 244, 247, 270,

271

and intuitionism, 47

for conditionalisation, 68

—E—

Elliptical references

to evidence, 1

Elliptical terms, 142

Ellsburg paradox, 130

Equivalence Analysis, 55–57, 65, 77, 225, 288

Event defined, 88

Evidence, 1

implicit, 151–53

sentences, 21

uncertain, 211–13

Exchangeability, 32–34

Exclusive disjunction, 63

Expressivism, 27–32

—F—

Families of predicates, 18

Families of probability functions. *See* Models

for degrees of belief; Many Models

FH conditionalisation, 106

Figures, 107–8, 227–30

Finkish. *See* Dispositions

Frege point, 29

—G—

Grue, 20, 241

—H—

Horizontal-vertical problem, 9

—I—

Implicature, 143

Independence, 32

Infinite utility, 39

Infinitesimals, 158

Internalism, 138

Intuitionist probability, 40

Intuitionist Probability. *See* Constructivist

probability

—K—

Kripke trees, 250–60

interpretation, 258

—L—

Liquidity, 326

and uncertainty, 330

motives, 329

Locke's Principle, 147

—M—

Marginal utility
 of money, 35, 48
 Mass function, 86
 MHP. *See* Monte Hall Problem
 Missing markets, 327
 Models for degrees of belief
 constructivist, 267
 imprecise
 Fagin-Halpern, 92–94
 Many Models, 74–80, 88
 Many-Urn, 94
 Single Model, 81
 updating, 100–110
 precise, 59
 updating precise models, 65
 Money, 321
 essential properties of, 325
 Monotonicity, 88
 Monte Hall Problem, 273, 277, 285, 294–96,
 339

—N—

New Principle, 159–65
 Non-projectible predicates, 19
 NP. *See* New Principle

—O—

Occam's Razor, 143

—P—

Paradoxes of Indifference, 19
 Pierino, 46–47
 Pignistic Probability, 291
 Platonism, 204
 Possibility spaces, 92
 PP. *See* Principal Principle
 Precise degrees of belief, 2
 Precisifications. *See* Supervaluations
 Preference
 axioms for rational preferences, 40
 Preferences
 axioms for rational, 130

independence, 130
 Reduction, 132
 transitivity, 130
 Principal Principle, 136–38, 159–65, 210
 defined, 136
 Principle of charity, 141
 Principle of Indifference, 57, 72, 195, 203, 209,
 243, 303
 Probabilistic Disjunctive Syllogism, 119
 Probability
 axioms, 4
 calculus, 3–6, 208
 conditional, 4
 interval-valued, 226–27
 Probability sentences, 3, 169–73
 comparative, 239
 relational, 134–38
 truth-aptness, 172
 Probability₂, 136
 Propositional gambles, 111
 Propositions
de nunc, 138
 simple, 206
 subject of probability sentences, 3, 134–36

—R—

Reduction, 288. *See* Preferences
 Reference class, 8
 Reflection, 46, 50, 65, 66, 67, 119, 226
 Relationism, 154
 Relativism, 153–56
 Risk aversion, 91

—S—

Shafer functions, 85–89
 State of confidence, 323
 Supertruth, 169
 Supervaluations, 166–94
 and first-order truths, 188
 and models of language, 185–93
 and validity, 182–85
 global validity, 183
 local and general, 193

local validity, 182
 moderate, 182
 penumbral connections, 168, 188
 permutation problems, 189
 precisifications, 168
 scope ambiguities, 173–82
 strong, 182
 weak. *See* orthodox

—T—

Three Prisoners Problem, 83–84, 102–3
 TPP. *See* Three Prisoners Problem
 Truth, 179–82
 T-schema, 167, 173–82
 Two-envelope paradox, 116–19

—U—

Uncertainty

and money, 321, 324–28
 defined, 307
 epistemic not aleatory, 319
 Keynes's definition, 306
 Uncertainty aversion, 89–92
 Undermining, 161
 Unemployment, 318
 Use-value and exchange-value, 43

—V—

Vacuous epistemic state, 93, 114
 Verificationism, 14, 243

—W—

Weight of arguments, 300
 analysed, 304
 and irrelevant evidence, 302