Universiteit Utrecht

Master's Thesis

The Corruption of Epistemic Puritanist Bayesianism

Author:
Jacco Oosterhuis
Student ID: 3714225

Supervisors:
dr. Gregory Wheeler
dr. Janneke van Lith

A thesis submitted in fulfilment of the requirements
for the degree of Master of Philosophy

in the

Department of Philosophy and Religious Studies

August 31, 2015

“Es fällt dem Philosophen nicht leicht, auf die Erkenntnis zu verzichten, aber es ist wahrscheinlich die große werdende Erkenntnis des zwanzigsten Jahrhunderts, daß man es tun muß.”

Robert Musil

Introduction

What is probability? As a mathematical question, there is almost complete consensus: it is the mathematical theory built on the axioms of Kolmogorov. As a philosophical question, the analysis of the *meaning* of probability, there are many different interpretations: probability can be frequency of similar events, a propensity to in the long run show an amount of specific outcomes, a degree of implication, or a degree of confidence of an agent that something is the case. This last theory is subjective probability theory, built on the famous thesis of De Finetti that "PROBABILITY DOES NOT EXIST"(de Finetti 1974, x). That is, objective probability does not exist. According to the subjectivist there is only uncertainty of agents about what is the case and mathematical probability theory represents this uncertainty; to say that some event has a certain probability always means that that event has that probability *for him or her. An agent* considers something to be so-and-so likely. This subjective meaning of probability does not necessarily exclude other interpretations—one can try combining multiple philosophical interpretations of probability and argue for pluralism with respect to the concept of probability—but the subjective interpretation is a very celebrated account of probability and very influential.

But it is not enough to say that probability is a representation of uncertainty; one has to say what this uncertainty is. A widespread understanding of subjective probability theory is that it is about rational *degrees of belief*: beliefs come in degrees, and if one follows certain rationality conditions in one's degrees of belief, one's degrees of belief are probabilities. Better put, Kolmogorov's axioms are representations of these rational degrees of belief. Depending on what one means by "degree of belief", different accounts of what rationality amounts to can be given; the most applauded is that given by the *Dutch Book argument*, which says that if you are incoherent in your degrees of belief, you will be open to a series of *bets* based on those degrees of belief which will guarantee sure loss. The Dutch book argument was first constructed by Ramsey (Ramsey 1926) and de Finetti (de Finetti 1972,

1980) and is based on the betting interpretation of degrees of belief, where a degree of belief is defined as a price for betting contracts; the Dutch Book argument is that any agent had better not have degrees of belief that are incoherent, because then she is open to a series of bets such that she would lose no matter what would happen—that would be irrational. Another similar interpretation of subjective probability, also by de Finetti (1974), defines subjective probability on the basis of forecasts, which are penalized according to how close the forecast is to the truth. This penalty is given by a scoring rule over forecasts, where this penalty is in some measure of linear utility, and incoherent forecasts would suffer an unnecessarily high penalty—posting incoherent forecasts is irrational. A third also very popular account of subjective probability theory is based on representation theorems of comparisons of likelihood of propositions; representation theorems are mathematical results that establish a representation given that certain conditions are satisfied, in this case certain coherence conditions are satisfied by comparative judgements of likelihood.

It is a small and tempting step from taking subjective probability theory as a theory of rational degrees of belief to taking it as an epistemological framework: the latter step would provide for an entirely different, and more fine-grained, framework for beliefs and epistemology. Furthermore, it would cash out the normativity of probability theory as rationality conditions of beliefs—a desirable result given the problems of establishing the normativity of logic over traditional beliefs. This interpretation of subjective probability as an epistemological theory is central to *Bayesian Epistemology*. However, the understanding of degrees of belief as *betting prices*, or as forecasts that are penalized, has left many epistemologists unsatisfied: even if agents are sometimes disposed to bet on events being true according to their degree of confidence in events or propositions, this is not generally the case. They state that the presence of subjective values in both the betting and forecasting interpretation has nothing to do with beliefs and epistemology. They want a *purely epistemic* subjective probability theory and a purely epistemic explanation of the rationality conditions; an agent's personal values and preferences, and the penalties in both the scoring rule and betting interpretations in measures of non-epistemic values, should not and do not matter for the *purely epistemic* value of degrees of belief. This *Epistemic Puritanism* has become a very popular camp in Bayesian epistemology.

However, there is more than one problem for such a purely epistemic construal of subjective probability theory. The first is that of a definition or precise characterization of degrees of belief: what *is* a degree of belief and how are degrees ascribed to beliefs? On the classical account of de Finetti, degrees of belief are defined as betting prices or forecasts, where beliefs become numerical through a measurement scenario. But this definition through one of the measurement scenarios of de Finetti is rejected by the Epistemic Puritanists. Yet no alternative is put in its place; no measurement device is provided by Epistemic Puritanists which allows us to ascribe degrees to belief in a nonambiguous way and to make sense of the notion of purely epistemic degrees of belief. The second problem is to specify what is required for a probability representation of degrees of belief and why an agent should conform to the requirements for such a probability representation; why it is epistemically rational to be coherent in one's degrees of belief, and what "epistemically rational" means. If a purely epistemic interpretation of subjective probability theory does not establish that rational degrees of belief can be represented by a probability function, it is not a theory of subjective probability.

I will here consider the hopes and plausibility of construing subjective probability theory as about purely epistemic degrees of belief that are independent of values and preferences. The research question overarching this entire work is: is epistemic subjective probability theory conceptually plausible? This requires investigating the following sub-questions:

- What are purely epistemic degrees of belief and (how) can we ascribe numbers to purely epistemic beliefs?
- What conditions can be justified from the purely epistemic standpoint on these degrees of belief, such that they are (represented by) probabilities?
- Can any of the classical accounts of subjective probability theory be adapted so as to be a purely epistemic theory of subjective probability?

I will answer these questions by starting from the notion of degree of belief and Epistemic Puritanism, and then analyzing different attempts at constructing a purely epistemic subjective probability theory.

Though attractive from the standpoint of traditional epistemology, I will argue that this purely epistemic subjective probability theory is seriously flawed. The main problem is that the successful and influential account of de Finetti depends essentially on values, both in the definition of degrees of belief and in the justification for the coherence conditions on these degrees of belief. Another approach, that of a much weaker notion of degrees of belief based on qualitative considerations of likelihood, depends on very strong axioms which are difficult to justify from the purely epistemic standpoint. As such, the epistemic approach fails both to supply a successful and strong enough notion of degrees of belief and also to justify rationality conditions such that subjective probability theory goes through. The thesis that I shall here defend is that purely epistemic subjective probability theory is conceptually implausible. Adapting de Finetti's (1974, x) well-quoted phrase, my thesis (perhaps paradoxically, and a little provocatively, but nonetheless genuinely,) is simply this: purely epistemic subjective probabilities do not exist; but this is not a problem.

It is not a problem because epistemic puritanism is ill-motivated. The criticisms about the traditional account of de Finetti and Ramsey are very often misguided, because they arise from a misunderstanding of de Finetti and Ramsey's theories. Firstly, de Finetti and Ramsey's theories are not based on (purely epistemic) degrees of belief, but are based on *dispositions to act*. Secondly, many criticisms focus on problems with regard to the betting interpretation and the Dutch Book argument; problems de Finetti was well aware of and which led him to prefer the forecasting interpretation that uses scoring rules over the betting interpretation.

This thesis is structured as follows. Chapter 1 will involve a lot of stage-setting and here I will introduce and analyze the essential notion of degrees of belief. In this chapter too we will try to obtain a probability representation of a very weak notion of degree of belief: a notion of degree of belief that is only qualitative, not quantitative, namely comparative judgements of likelihood. This is done through representation theorems, which are results which we will also discuss extensively in the end of chapter 1. Although this approach is possible and popular, the conditions required for a probability representation are very difficult to justify from the purely epistemic perspective. In chapter 2 and 3 we will consider the very influential approach of de Finetti, where chapter 2 will consist mainly of explanation which we will then critically discuss in chapter 3. In this latter chapter we will try to take the setup of de Finetti and interpret it in a purely epistemic way, where degrees are not taken to be dispositions to act. As we shall see, this purely epistemic path is blocked, for the essential elements which allow for a measurement scale to assign numbers in de Finetti's setup are missing in the purely epistemic theory. But even supposing that there are degrees of belief in a purely epistemic sense, as we shall do in chapter 4, subjective probability theory is still a no-go: a purely epistemic justification of the coherence conditions on degrees of belief fails. Two recent attempts at such a purely epistemic justification, one that adapts the betting scenario and one that adapts the scoring rule scenario, both do not succeed. Especially the latter attempt has been very popular in recent debate, even though it fails in its foundations: it requires a notion of epistemic utility which is conceptually very problematic. Lastly, in chapter 5 we will summarize the results of this thesis and briefly discuss the relation of the results with other research in epistemology.

Contents

Chapter 1

The qualitative approach

Subjective probability theory is a theory based on a very simple idea: probability is rational subjective uncertainty. This subjective uncertainty is uncertainty of any agent about what is the case and can be, if subject to certain rationality conditions, represented by a probability function. Probability functions are mathematical functions defined by Kolmogorov's axioms. But what this subjective uncertainty is, is unclear from just this characterization: what is being represented by probability theory? Once we have an answer to that question, the next matter is what it is for that what is being represented to be rational. This rationality is cashed out by certain axioms or coherence conditions on beliefs, acts, or opinions, where not following these coherence conditions is irrational in some way. What the coherence conditions are depends on how we construe subjective probability theory.

An important matter that must first be settled is whether subjective probability theory is taken to be normative or descriptive: whether the theory says that the subjective uncertainty *should* be representable and whether agents *should* follow coherence conditions and axioms, or whether it says that subjective uncertainty *is* representable and agents *do* follow coherence conditions and axioms. A descriptive interpretation would be a psychological or sociological theory, which is not what we are interested in here. The normative interpretation is the one adopted in this thesis throughout: subjective probability theory tells us how agents *ought* to act or reason in situations of uncertainty. But it is important to note that all normative claims in this work are normative *about* something which the normative theory posits to exist. So if the subjective probability theory is taken to be prescriptive of degrees of belief, it assumes or requires that there *are* degrees of belief and that these should follow certain axioms; it requires beliefs to have a certain specific structure to get off the ground.

1.1 Subjective probability theory as epistemology

A popular and widespread conception of subjective probability theory is to see it as an epistemological framework and theory. Whereas work in traditional epistemology primarily concerns the characterization of knowledge, of which the classic answer (going back to Plato) is that knowledge is justified true belief, the conception of subjective probability theory as an epistemological theory rejects the primacy of the concept of knowledge and instead analyses degrees of belief, evidential relations, and likelihood; this conception can be called *Bayesian Epistemology*. Bayesian Epistemology is a very broad term and the position is far from a single clearly unified movement; for example, the "epistemology" need not strictly be like traditional epistemology but can only be marginally like it, and not be concerned with characterizing the notion of knowledge at all. Also there are many different positions in Bayesian Epistemology on what updating principle should be adopted—how degrees of belief ought to change over time.

Bayesian Epistemology usually consists of three aspects (Easwaran 2011, 312): beliefs come in degrees; these degrees of belief ought to satisfy certain rationality conditions which correspond to the axioms of probability theory; degrees of belief ought to be updated in a certain way—conditionalization is the most popular and famous updating rule. The second of these three topics is a *synchronic* topic and concerns the rationality conditions of degrees of belief *at* a certain time. The third topic is a *diachronic* one, about the rationality conditions of degrees of belief *over* time. However, since the synchronic part is already big enough to fill entire books, we will throughout this work only be concerned with the synchronic part and will not discuss any diachronic conditions. Thus, we will restrict ourselves to the first two aspects of Bayesian Epistemology. It is also important to note that the third topic and many other topics within Bayesian Epistemology require, as a foundation, a synchronic theory to build upon; if the part we will discuss in this thesis is unsound in a certain construal of Bayesian Epistemology, the other projects in that construal also fail as a result.[1]

1. There are many further topics related to subjective probability theory as an epistemology: for example, the relation between degrees of belief and traditional beliefs, and the relation with Bayesian Epistemology and the concept of knowledge. One popular approach here is to postulate the existence of so-called "bridge laws" which say when having specific degrees of belief counts as believing some proposition. See (Huber

1.1.1 Epistemic Puritanism

One popular construal of Bayesian Epistemology is to try to say that subjective probability theory is a theory purely and only about beliefs and that the justification for rationality conditions on beliefs can be given a purely epistemic foundation. Beliefs on this view are attitudes whose primary aim is the truth: beliefs can and should be evaluated according to their truth, not according to pragmatic virtues and values. On this approach, probability theory is a theory about purely epistemic rational beliefs. We will call this interpretation of Bayesian Epistemology **Epistemic Puritanism**. The idea is that subjective probability is a representation of degrees of confidence in propositions, or degrees of belief, and that these degrees of belief are psychological attitudes very much like beliefs. They are "pure" in the sense that they can be evaluated on purely epistemic considerations and are not essentially connected to values.[2] The purely epistemic approach says that the rationality conditions on degrees of belief can be argued for by considering only purely epistemic values; whatever values and preferences the agent has is irrelevant for the purely epistemic construal of subjective probability theory.

A purely epistemic approach to Bayesian Epistemology is, as described, a theory of rational belief: it says how any agent's beliefs *ought* to be structured. Besides the rationality conditions, which we will construe as a structure which beliefs ought to satisfy, it therefore requires a notion of belief strong enough to make specific rationality conditions tenable. Understanding the notion of belief, and the way (Bayesian) epistemology relates to it is therefore very important.

1.1.2 Folk Psychology and beliefs

Traditional or "full" beliefs are attitudes or states of what is called "folk psychology": they are the attitudes we attribute to other people in everyday life, primarily when explaining their actions, along with desires, intentions, thoughts, and the like. This folk psychology is

and Schmidt-Petri 2009) for a collection of papers on this topic. We will not discuss these topics further here.

2. We will discuss many proponents of this position throughout this work. Some important authors are David Christensen (2004), James Joyce (1998, 2009), Richard Pettigrew (2013) and Hannes Leitgeb (2010a), and Hillary Greaves and David Wallace (2006).

the theory[3], stance[4] or framework we use when explaining primarily behaviour; when we approach humans as agents that make intentional actions. Indeed, it can been said that the primary purpose of folk psychology is the explanation and prediction of behaviour and that in explaining behaviour as rational, we cannot get around the intentional—get around beliefs, intentions, and desires.

The attitudes of folk psychology are primarily attributed on the basis of behaviour, but beliefs and other psychological attitudes are not taken to be definable as behaviour: they are postulated as underlying, or perhaps causing, behaviour. Despite the central role of beliefs in leading to action, not all beliefs lead to action and beliefs cannot be strictly identified with actions, nor can one read off, from one's actions, what someone believes or believed. Behaviour can be caused (in a colloquial sense of causation) by many different beliefs and desires: in making a cup of coffee I can either want caffeine, want to stay awake, want the taste, or simply want to make a cup of coffee; the belief that caused my action could be that making coffee satisfies any one of these desires, and a lot more besides. Furthermore, explanations in folk psychology are made, essentially, by referring to an agent's beliefs and desires—"why is he making coffee?" "Because he feels like it", or "because he believes it will do him good"—so that any reduction to or identification with behaviour of beliefs and desires becomes impossible within folk psychology. We will revisit these points later in chapter 3.

The question to the nature of full beliefs is important for this thesis because the answer also relates to degrees of belief: if beliefs are said to come in degrees, it seems desirable that these degrees of belief are similar in many ways to full or traditional beliefs—otherwise they had better not be called degrees of *belief*. In philosophy of mind, beliefs, intentions, desires, and the like, have generally been characterized as representational of character, meaning, roughly, that they are about something; beliefs represent aspects of reality or are about reality (Schwitzgebel 2015). This is reflected in the characterization of beliefs as specific propositional attitudes.[5] But this representational character is not all there is

3. For this conception see especially (Sellars 1956). Whether it is a genuine theory in the way Sellars describes is disputable; it is not clear that folk psychology involves (or can involve) laws. See also (Millikan 1993, 52–54; Davidson 1980b, 1980c).

4. See (Dennett 2002).

5. It differs from other propositional attitudes in its direction of fit towards propositions: beliefs are adopted according to the way the world is, while intentions try to change the world in some way. Actions also do not have the same direction of fit as beliefs: actions are about changing reality, whereas beliefs are changed in accordance with reality. See (Anscombe 1957). This aspect is also emphasized by Levinstein

to beliefs; the representational character can be taken as the primary characteristic or just some element of a wider picture. As to the nature of beliefs and their *identity*, positions differ strongly: some emphasize the representational character of beliefs, others its functional role in action and behaviour, still others the importance of interpretation of agents in social contexts (Schwitzgebel 2015). Regardless of what philosophical characterization of beliefs we choose, any such theory of belief is a theory of the *identification* of beliefs. For a functionalist theory, the essence of beliefs is its function in leading to behaviour or its function in an organism, and beliefs are identified according to their (normal) functions; for representationalists, beliefs are identified according to what they represent; for interpretationists, beliefs are identified according to interpreting the behaviour of an agent and treating her as an intentional being. And for a traditional approach beliefs are identified according to introspection—beliefs are what they seem and seem what they are. This will be important shortly, when trying to characterize (and identify) degrees of belief, which we will do after considering beliefs and their role in theories of knowledge.

Besides philosophical theories of rationality and action, where the common sense concept of belief of folk psychology is prevalent, beliefs have been central in many other areas of philosophy, most importantly epistemology. Traditional epistemology is primarily interested in the characterization of knowledge, of which the classic answer (going back to Plato) is that knowledge is justified true belief. Traditional epistemology is also normative: the *aim* of belief is frequently taken to be knowledge and beliefs are, for that reason, evaluated according to the truth of propositions. This is the *truth-norm* of traditional epistemology. Traditional epistemology takes belief as a folk-psychological primitive and tries to characterize necessary and sufficient ingredients for knowledge that beliefs have to satisfy.

1.1.3 From beliefs to degrees of belief

With respect to Bayesian Epistemology and Epistemic Puritanism, the previous section gives us an idea of how epistemology and subjective probability theory might coincide. The idea is that subjective probability theory is a theory that applies to the folk-psychological attitudes called beliefs. Traditional “full”, all-or-nothing, beliefs are not fine-grained enough to allow for a probability representation. In traditional epistemology, one either

and Konek in relation to degrees of belief and acts, whose position we will discuss in chapter 4 (Levinstein and Konek 2015). This characterization of beliefs as propositional attitudes is not universally shared, but its rejection does no damage to the points made in this thesis.

believes or does not believe something.[6] Being non-numerical and very rough of structure, beliefs as just folk-psychological attitudes do not allow for a probability representation. The classical step of subjective probability theory is therefore to say that beliefs come in degrees: we do not just believe propositions to be false or true, but we frequently or even normally believe propositions to be more or less *likely* and have a higher or lower degree of confidence in propositions. So it can be said that we believe something to a certain *degree*. For example, I considered it unlikely that it was going to snow in April 2015 in Munich, but I was not sure of the matter. (It did, thanks to the cyclone Niklas.) In any case, I was still more confident in that proposition than in the proposition that I could write a master's thesis on Epistemic Puritanism in one month—which I could not. Hence, my degree of belief in the former proposition was still higher than my degree of belief in the latter, although quite low.

But talking of degrees of belief here is speaking very roughly; it is quite a step further to hold that beliefs come in numerical degrees. We could of course attach and distribute numbers to our own beliefs willy-nilly, but this would be arbitrary and ad hoc—surely not precise enough for any philosophical theory. I cannot say I believed it to a precise degree of 0.268—I would not know how to ascribe numbers, nor what the numbers meant—but I did consider new snowfall in April 2015 pretty unlikely. Furthermore, such an introspective account is now also generally rejected with respect to traditional full beliefs: we do not have introspective direct and infallible access to what we believe.[7] To reinstall this old rejected picture in our new theory of degrees of belief without justification can hardly be called a step forward, nor can it be said to be plausible.

Having rejected the introspective account, we can perhaps hope to provide more meaning to the notion of degree of belief by looking at the other theories of identification of beliefs, discussed above. What we need in this case is a criteria that ascribes a precise number according to some precise procedure—we do not want ad hoc and arbitrary num-

6. Under "not believing" I take both believing the negation of some proposition and suspension of judgement. This last possibility, to suspend judgement on whether some proposition is true or not, might be said to not be captured fully by subjective probability theory: Bayesian epistemology always says that one believes something to some degree, and to suspend judgement about some proposition is to say that one does not believe something to be true or false; it therefore presumably also says that one does not believe something to be true or false to some degree. How suspension of judgement relates to subjective probability theory is a difficult question related also to the earlier-mentioned bridge-laws between degrees of belief and full beliefs, which will not be discussed further here.

7. See (Davidson 1987; Burge 1986) for two important papers on this topic.

bers. Unfortunately, I do not see how the traditional conceptions can give precise meaning and unequivocal ascription of numbers to degrees of belief. Consider representationalism, where belief states are individuated according to what they represent. Such a representation relation does not offer a clear way of ascribing numbers. My believing snow in Munich in April 2015 unlikely is a belief about the weather in Munich in April 2015, but that does not tell us by itself how likely precisely I considered it. Nothing in the representation relation seems to offer a way of ascribing numbers. Or consider functionalism, where a belief state is individuated according to its function performed in causing behaviour (in normal circumstances), can we use this to ascribe numbers to beliefs? The idea would then presumably be that a degree of belief is numerically determined by its function. If the function of belief is to lead to action, this does lead to a way of ascribing numbers, namely on the basis of action, but then ascribing numbers to beliefs can be done *only* insofar as beliefs lead to actions; beliefs that are not related to actions (a possibility, even for functionalists) cannot be ascribed numbers. Otherwise, if we take the biological or anthropological function of belief to be the function of belief, I do not see how this could work. Interpretationalism seems even less up for the task.

The problem is that the only way of clearly assigning numbers is by some sort of measurement procedure. We need, therefore, a way of measuring beliefs, and the traditional accounts of the notion of belief do not clearly supply this, nor is it clear how folk-psychological beliefs are to be measured, according to the theories of beliefs just mentioned. A promising proposal here is to measure beliefs through actions; this is the idea that Ramsey explicitely used and the way de Finetti can also be interpreted as using, which we shall discuss in the chapters to come. However, we have already noted that beliefs cannot be identified with behaviour or actions for multiple reasons, and the same applies to degrees of belief: the connection between beliefs and behaviour is too weak to allow us to measure (degrees of) belief through behaviour or actions. We cannot even tell in straightforward cases what one thinks or believes, for sure: if one says “I believe that the cat is on the mat” he can still mean and believe any number of things, depending on the context and other matters.

Another possible avenue is to say that the ascription of precise numerical degrees of belief is a strong idealization, but that actual degrees of belief are vague, which is then explained as beliefs being interval-valued. I consider something likely to a degree of 0.6–0.85, say. I think this is, purely as such, just as arbitrary: it is still assigning numbers that are not unproblematically present or meaningful in the case of beliefs. Compare this with the belief that a table is somewhere between a meter and a meter and a half long: this is indeed uncertainty about the actual length and the interval is helpful, but we have here a clear way of assigning numbers and a clear *meaning* of these numbers. We will revisit and discuss both suggestions of the last two paragraphs in chapter 3.

These considerations at least give reason to believe that beliefs, as folk-psychological attitudes, do not come in numerical degrees. There are multiple starting points from which we can hope to attain a less problematic notion of degree of belief. One is to insist that any degree of belief is in principle measurable, and that this measuring scenario ensures specific numerical degrees. This is the approach of both Ramsey and de Finetti (Ramsey 1926; de Finetti 1974, 77–76) which we will discuss in the chapters to come. Let us here start with a notion of degrees of belief as *considerations of likelihood* of propositions; these are non-numerical and provide a quite plausible and intuitive structure of beliefs. These comparative judgements are meant to capture the intuitive meaning of "more likely than" (and "less likely as") in the following way: if I believe some proposition but not its negation, I consider it more likely than its negation. If I consider a proposition very likely, this can be taken to mean that I consider that proposition more likely than many other propositions.[8] We do not assume a numerical value of likelihood, but just a qualitative ordering among propositions: for this picture we only need to assume that beliefs are structured such that people can compare different propositions by their likelihood.[9] We understand the comparative judgements of likelihood as ordinal scale comparisons, not cardinal. What we want, in the end, is a probability representation of this likelihood relation.

1.1.4 Representation theorems

Such a probability representation can be ensured by a representation theorem. Representation theorems are mathematical proofs that establish that a certain mathematical structure can be taken as a mathematical representation of another (usually non-numerical) structure. This is so only if the latter structure satisfies certain axioms. As an example, take the case of comparative judgements of likelihood: the structure we have here (which we will introduce shortly) has a relation which is non-numerical and only an ordinal scale; this relation can be made precise and numerically measurable by a mathematical repre-

8. This last point depends on the underlying algebra and other properties of the structure. It is meant as a rough claim for explanatory purposes.

9. It is important to note here that by "likelihood" I mean a purely subjective consideration, *not* an objective notion such as chance, for which the term is also sometimes used. Likelihood must also be distinguished from probability, by which I will always mean a mathematical probability *function*.

sentation. With only an ordinal relation, we cannot say how much more likely something is or quantify how likely we consider something; on the other hand, with a mathematical representation we can clearly express this by numbers, which gives us a clearer and more informative structure.

Such a representation is a homomorphism between structures: a mapping from one relational structure to another, which preserves the structure of the former. Sometimes such a representation is also an isomorphism, meaning that the mapping admits an inverse, but this is not generally the case and also not for the cases we will discuss here. A representation theorem is a proof that such a homomorphism of a relational structure into a numerical structure exists; such a theorem is nontrivial and an important mathematical result, for only if certain axioms are satisfied by the relational structure do we have such a homomorphism. So, on the approach by means of representation theorems of subjective probability theory, only if we have a representation theorem does the theory as a whole succeed; otherwise probability is not a representation of some subjective qualitative relation.[10]

1.2 Qualitative probability: considerations of likelihood

In our case, we want a quantitative representation of our qualitative likelihood relation. We can start with the binary operation $\succcurlyeq$ among propositions, where $A \succcurlyeq B$ says that some proposition A is *at least as likely as* a proposition B.[11] The question is what conditions must be satisfied by this relation for there to be a numerical representation of it.

For this we need to introduce some formal setup. The starting point of qualitative likelihood is a triple $\langle \Omega, \mathcal{A}, \succcurlyeq \rangle$, where Ω is an outcome space $\{w_1, w_2, ...\}$—possible outcomes of some experiment or any state of the world the agent is unsure of. This sample space can be either countably infinite or finite. Ω constitutes a partition: one and only one w_i must occur. Every such element of Ω is therefore a state that is possible for the agent. $\mathcal{A}$ consists of subsets of Ω, and we will call the elements of $\mathcal{A}$ events or propositions[12] and denote these by $A, B, C, ...$ This family of subsets must be an algebra, meaning that it is nonempty,

10. In this section and the next I have drawn extensively from the treatment of the topic in (Suppes 2002) and (Krantz et al. 1971).

11. I here depart from the custom of denoting propositions with $p, q, r, ...$, and shall use instead $A, B, C, ...$ This step is made less awkward by the identification of propositions with subsets of the sample space.

12. Throughout this thesis we shall consider events and propositions to be identical: both are subsets of the sample space. We will use the two terms almost interchangeably. A proposition can be understood for

closed under union and closed under complementation. Lastly, the qualitative relation $\succcurlyeq$ is the relation "is at least as likely as" defined on $\mathcal{A}$. Two events being considered just as likely is indicated with $\approx$ and is equivalent to both $\succcurlyeq$ and $\preccurlyeq$ holding between the two events. Also, $A \succ B$ is equivalent to $A \succcurlyeq B$ and not $A \approx B$.

Although much of the literature uses this set-theoretical setup, de Finetti, DeGroot and others have preferred to use random variables instead.[13] Random variables are variables of which the actual value is unknown for some agent, and these random variables are denoted by $X, Y, \ldots$. Random variables can take any number of values or points within some *space* of possible values—de Finetti calls these "atomic" events (de Finetti 1974, 32). This space corresponds to Ω and random variables, as such, can be understood as functions from Ω to $\mathbb{R}$. The true 'point'[14], the actual value of the random variable, is a random point or state w in Ω, and $X(w)$ is the value X takes if w is the true point. Random variables are preferred by some because they allow for a general theory of *prevision*, of which probability is a specific case, and because of the ease of introducing new random variables rather than changing Ω (Walley 1991, 57–58). For our purposes, defining probability on random variables or on Ω can be considered equivalent, and random variables are used primarily in explaining the setup of de Finetti in the next chapter.

Events and propositions are special cases of random variables, namely indicator functions of subsets of Ω. An indicator function $\mathbf{1}$ of an event or proposition $A \in \mathcal{A}$ is a random variable which can only take the values 0 and 1, which yields a 1 if the event A occurs (if the proposition that A is true), and 0 otherwise. Thus, $\mathbf{1}_A$ yields a 1 if the true point or state w is an element of A. Thus events or propositions A and B are random variables that yield the value 1 if the true point w is in a subset A of Ω, and 0 otherwise.

A probability function $\mathbf{P}$ is a function from $\mathcal{A}$ to $\mathbb{R}$, which obeys the following conditions for every $A, B \in \mathcal{A}$:

1. $\mathbf{P}(A) \geq 0$.
2. $\mathbf{P}(\Omega) = 1$.
3. If $A \cap B = \emptyset$, then $\mathbf{P}(A \cup B) = \mathbf{P}(A) + \mathbf{P}(B)$.

our purposes as saying that some event is the case. We will understand "event" in a very wide sense: the moon being made out of Roquefort is also an event.

13. The notation used here is mostly that used by de Finetti (de Finetti 1974).

14. What these possible points are we will discuss in chapter 2.

These correspond to the axioms of Kolmogorov's probability calculus. As mentioned, for subjective probability we want the probability space $\langle \Omega, \mathcal{A}, \mathbf{P} \rangle$ to be a numerical representation of $\langle \Omega, \mathcal{A}, \succcurlyeq \rangle$. In the case of random variables, the function $\mathbf{P}$ is only a probability function if X is an indicator function; if X is another random variable it is a *prevision* function. A prevision function is a function $\mathbf{P} : X \to \mathbb{R}$, for all random variables X. It corresponds to a (coherent) expectation of a random variable and will be important in chapter 2 where we will discuss de Finetti's approach. Probability is a particular case of prevision, namely an expectation that a random quantity takes a certain value; that a random quantity takes a certain value, is a proposition or event.

Aside from saying that the probability function is a mathematical representation of the qualitative structure defined by $\succcurlyeq$, we can also say that $\mathbf{P}$ *agrees* with $\succcurlyeq$ which is equivalent and frequently used in the literature. $\mathbf{P}$ agrees with $\succcurlyeq$ if, for all $A, B \in \mathcal{A}$,

$$A \succcurlyeq B \Leftrightarrow \mathbf{P}(A) \geq \mathbf{P}(B)$$

It is not difficult to realize that not every qualitative consideration of likelihood will have a probability representation. If the considerations of likelihood are incomplete, so that not for every A and B either $A \succcurlyeq B$ or $A \preccurlyeq B$, there will not be a unique probability representation but a set of probability functions which agree with the (incomplete) ordinal structure; if the ordering is cyclic, so that for some A, B, C, $A \succcurlyeq B$, $B \succcurlyeq C$, and $A \preccurlyeq C$, any numerical representation will fail. As will become clear, to ensure such a representation quite strong conditions on the structure defined by $\succcurlyeq$ must be imposed. A good starting point and important first attempt for sufficient conditions on the structure were famously proposed by de Finetti (de Finetti 1980), where for any $A, B, C \in \mathcal{A}$:

1) $\succcurlyeq$ is a weak ordering on $\mathcal{A}$, meaning that for any two events A and B, $A \succcurlyeq B$ or $B \succcurlyeq A$ holds, and that if $A \succcurlyeq B$ and $B \succcurlyeq C$, then $A \succcurlyeq C$ (transitivity) ;

2) For any A in $\mathcal{A}$, $A \succcurlyeq \emptyset$ holds ;

3) $\Omega \succ \emptyset$;

4) If $A \cap C = \emptyset$ and $B \cap C = \emptyset$, then $A \succcurlyeq B$ if and only if $A \cup C \succcurlyeq B \cup C$. (This says that the relation is closed under logical conjunction.)

De Finetti conjectured that these conditions were adequate to ensure a numerical representation of qualitative probability. It was shown in (Kraft, Pratt, and Seidenberg 1959) that this is not the case, for infinite nor finite Ω. Using a structure where Ω contains five elements, Kraft, Pratt, and Seidenberg constructed a model which complies with the properties just listed, but where trying to obtain a probability representation leaves us with the result that $0 > 0$ and thus inconsistency; we cannot have a numerical representation of this model.[15]

De Finetti proved early on (de Finetti 1980) (much earlier than the result of Kraft, Pratt, and Seidenberg) that a probability representation can be ensured if Ω can be divided into a partition of arbitrarily many events of equal probability:

5) For every positive integer n, there exists a partition $C_1, ..., C_n$ of Ω, such that for any two integers $i, j = 1, ..., n$, $C_i \in \mathcal{A}$ and $C_i \approx C_j$ (Krantz et al. 1971, 207).

Of course, the sample space is by no means always of such a character. This can be ensured, however, by multiplying any event with, say, an arbitrarily large number of tosses of a fair coin. Given that the amount of coin tosses is big enough, the sample space can be subdivided in the required way. Savage (1954) proposed a similar (but stronger) condition:

5a) For any two events A and B in $\mathcal{A}$, if $A \succ B$, there is a partition of Ω in events $C_1, ..., C_n$, such that $A \succ B \cup C_i$, for every $i = 1, ..., n$ (Savage 1954, 38).

This is also explained by Savage as that, if you consider A more likely than B, there is a sequence of coin tosses such that you would continue to consider A more likely than the union of B and any particular sequence of heads or tails (Savage 1954, 38–39). Here too it is required that Ω is infinite.

DeGroot formulated a similar condition in terms of auxiliary experiments, for which he used random variables. DeGroot's condition consists of the existence of a specific kind of random variable: an *uniformly distributed* random variable X, where X is such that $0 \leq X(w) \leq 1$, for every $w \in \Omega$. X is uniformly distributed over $[0, 1]$ if the probability that X will belong to any interval I with end points a and b (with $a \leq b$) is equal to the length of that interval. DeGroot proves that this is sufficient for a numerical representation of qualitative probability (DeGroot 1970, 76–78). This condition also comes down to Ω being divided into infinitely many events of equal probability. The existence of such a random variable can always be ensured by the existence of what DeGroot calls an "auxiliary experiment": an experiment where "the value of a random variable having the appropriate uniform distribution is observed"(DeGroot 1970, 76).

15. See (Fishburne 1986, 337) or (Suppes 2002, 228–229) for accessible explanations of the result.

This condition and those of De Finetti and Savage are plausible enough from the statistician's point of view, where the statistician wants to calculate the probabilities of his own considerations of likelihood. This was what DeGroot was interested in and it is successful in this regard, for here obtaining a probability representation is desirable from the agent's perspective itself. In general, however, a normative case must be made why such a probability representation is desirable, or why the axioms should be satisfied. No real normative case is given for these conditions. Also, a major downside of this condition is that the sample space is required to be infinite (Fishburne 1986, 341; Suppes and Zanotti 1976). The condition requires, as a basic example, that any consideration of likelihood can be combined with very, very large sequences of coin flips. The first problem with this is that coin flips have nothing to do with beliefs purely as such and why an agent should consider the problem in this way is unclear. To say that this is so that a probability representation can be obtained begs the question in an obvious way: we want to establish, after all, why an agent should want to have beliefs such that a probability representation is possible. Of course, given that the goal is to construct a probability theory or a decision theory, we want a probability representation, but this does not tell us why the considerations of likelihood should conform to the axioms; as we have already established, we want a normative theory, and the agent might very well be indifferent to being representable by a probability function or not. Following the axioms in one's comparative judgements of likelihood should be argued to be rational. Also, it is unclear that any agent *can* consider any problem in the way the axioms demand, precisely because they demand the sample space to be infinite, and it is not clear agents can reason in this way; if the requirement is normative, it had better follow the dictum "ought implies can".

Sufficient conditions that do not require an infinite sample space have also been established.[16] Kraft, Pratt, and Seidenberg already listed necessary and sufficient conditions to ensure a numerical representation of a probability function, but Dana Scott gave sufficient conditions that are much easier to understand (Scott 1964). Because Scott's axiom is a powerful axiom, we do not need the conditions 1), 2), 3) and 4) (nor of course the extra conditions of De Finetti or Savage) but can do with a shorter list (Suppes 2002, 229), for any A, B in $\mathcal{A}$:

1') $A \succcurlyeq B$ or $B \succcurlyeq A$;

16. Many of the conditions are sufficient, but stronger than strictly required. So is also Scott's axiom. For necessary and sufficient conditions see (Suppes and Zanotti 1976). For an overview of conditions see (Krantz et al. 1971, 202–208).

2') $A \succcurlyeq \emptyset$;

3') $\Omega \succ \emptyset$.

Scott's extra axiom is an algebraic condition on indicator functions of events:

4') If the sums of the indicator functions of two sequences $A_1, ..., A_n$ and $B_1, ..., B_n$ are equal, and $A_i \succcurlyeq B_i$ for $0 \leq i < n$, then $B_n \succcurlyeq A_n$.

This is a very strong assumption on qualitative probability to make. It says that any element of Ω belongs to the same number of true events A_j as B_j. If the number of truths (indicator functions that take the value 1) in two sequences are equal, then the likelihood relation should reflect this: the likelihood considerations should be balanced accordingly. If you consider $A_i \succcurlyeq B_i$ for every $0 \leq i \leq n$ in the respective sequences, but the indicator functions of the sequences are balanced, then $A_j \succ B_j$ should not hold for any number j.

Similar conditions sufficient *and* necessary for agreement with a (unique) **P** have been established by Suppes and Zanotti (1976) for both the infinite and finite case. This requires going from the algebra $\mathcal{A}$ to an algebra of what the authors call an algebra $\mathcal{A}^*$ of extended indicator functions. They explain this as follows: $\mathcal{A}^*$ contains the indicator function of any event A in $\mathcal{A}$, and for any two elements A^* and B^* of $\mathcal{A}^*$, $A^* + B^*$ is in $\mathcal{A}^*$. Thus, any element of $\mathcal{A}^*$ is an integer-valued function on Ω. To have $A^* \succcurlyeq B^*$ means that the expected value of A^* is at least as high as the expected value of B^*. We shall not here go through the exact construction given in (Suppes and Zanotti 1976; Suppes 2002, 230–232), but the important axiom for agreement with a probability measure is an axiom very similar to the Archimedean axiom of measurement theory (Krantz et al. 1971): "if $A^* \succ B^*$ then for every C^* and D^* in $[\mathcal{A}^*]$ there is a positive integer n such that $nA^* + C^* \succcurlyeq nB^* + D^*$" (Suppes and Zanotti 1976, 435). This condition is important for our thesis because of the formulation in terms of *expectations* of indicator functions. Considerations of likelihood are adapted into expectations of events.

These technical results are very significant, as a short recap will make clear. Not only must an agent's evaluations of likelihood of propositions be a weak ordering on all propositions on some possible state of affairs (a probability space), for that together with the other requirements of De Finetti is not enough to ensure a representation by a probability structure and thus not enough for subjective probability theory. The proposed conditions given by de Finetti alone might be said to be plausible to some extent (the fourth being the most contestable one), but even for these it is not immediately clear that these can be defended as applying to degrees of belief. Normatively—the way we are interested in for this thesis—the case for satisfying the conditions is difficult: *why* should one abide by the conditions? Assuming, as we have done, that subjective probability theory is a normative theory of (degrees of) beliefs, a normative case for the axioms must be established. Furthermore, as Kyburg has noted (Kyburg 2003), Keynes famously rejected the axiom that probability is a total ordering: some propositions are incomparable. If this is right, the path is immediately blocked.

One could perhaps say that violations of the axioms constructed for the structure defined by $\succcurlyeq$ is akin to violations of logic; one feels, intuitively, that there is something wrong with the violations. Close inspection, it might be said, reveals automatically that there is something intuitively wrong in not conforming to the conditions. This was Savage's view with regard to conditions of preferences (Savage 1954, 20–21), but this will hardly convince anyone who opposes the conditions: something I do see as a requirement of any theory of rationality. Of course, this intuitiveness can be said to be the case of traditional logic, too: I think Lewis Carroll's Achilles did not do a bad job in trying to prove modus ponens to the turtle (Carroll 1895). Someone who doubts the validity of basic operations of logic can perhaps not be convinced of their validity—we would not know what to say to this person—and a similar argument can perhaps be made for the axiom described here.[17] However, this step is implausible for the more technical *sufficient* conditions.

Scott's axiom is a hard condition on degrees of belief to argue for: it requires an *algebraic* condition of additivity of events. Scott himself noted the "unpleasant" nature of the condition being not a "strictly Boolean condition" (Scott 1964, 246). A rational case for Scott's axiom or an equivalent one is required. The problem is this: conditions on indicator functions are difficult to make sense of from the purely epistemic viewpoint. Of course, indicator functions are just propositions, but when considering the likelihood of propositions or events we do not *sum* propositions. We *can*, but then the normative question as to why we should always do so can still legitimately be posed. Similarly, it is implausible to say that when considering likelihood we consider expected values of propositions, as in the theory of Suppes and Zanotti. We do not, normally, have values in our head, as already said, which is how I understand this condition. So it seems reasonable to conclude that no real successful normative case can be made for any of the principles

17. See (Christensen 2004; Kolodny 2007; MacFarlane 2004) for some works on the normativity of traditional logic.

that secure a probability representation of only the considerations of likelihood.[18]

1.2.1 Measurement of likelihood and preferences

This is not to say that I consider all the axioms listed implausible in general. If the conditions are posed as conditions of *measurement*, any situation can be presented to an individual as a proposition where the agent is *asked* what her expectations are of some event. This would be a plausible way of construing, for example, the conditions given in (Suppes and Zanotti 1976) in terms of extended indicator functions. This is not to say that the agent has mental expectations in her head, but only that in a measurement scenario she would elicit such values.

Probability representations can not only be derived from the subjective considerations of likelihood, but can also be derived from an agent's preferences among acts or choices. If an agent's preferences over acts or choices satisfy certain axioms, a probability and *utility* representation can be derived from these preferences. Utility is a mathematical representation of an agent's values. Also, if specific well-defined options are given to the agent in which utility is known, an agent's preferences over these options can also be used to derive a probability representation. Such is the approach of Anscombe and Aumann, where through the (double) application of utility theory over preferences over certain kinds of lotteries a probability representation is established (Anscombe and Aumann 1963). Another approach is that of Savage (1954), where a utility and probability representation is derived from an agent's (rational) preferences in one go.[19] These results of decision theory have been very influential, but they are not too important for this thesis about epistemic puritanism because they are even more about *decisions* rather than about beliefs, than the theories of Ramsey and de Finetti. If the setup of De Finetti is not acceptable from the purely epistemic standpoint, neither are different setups based on decisions and preferences over acts.

18. James Joyce (Joyce 1998, 602) does think a normative case can be made, but this case depends fundamentally on his interpretation of scoring rules and accuracy. We will discuss Joyce's construction of epistemic probabilism extensively in chapter 4.2.

19. See also (Fishburne 1981) for a wide review of theories of decision making in situations of uncertainty.

1.3 What do representation theorems establish?

Even if the coherence conditions are made plausible, some philosophers (Meacham and Weisberg 2011) have said that representation theorems do not establish anything of strong normative import at all; they do not and cannot say anything about our actual preferences or degrees of belief. For a positive result of subjective probability theory, getting clear on representation theorems and the way they work—if they work—is vital.

Let us first summarize briefly what representation theorems are all about: given certain constraints on your considerations of likelihood, there is a representation of these by a probability function. We can say here that the agent has considerations of likelihood and that, if certain conditions are satisfied, she considers likelihoods of events *as if* she attaches probabilities to these. The conditions over comparative judgements of likelihood then establish that there is a representation with a unique probability function. In the case of obtaining a probability representation over preferences this uniqueness is not necessarily the case; utility functions are unique up to affine transformation.

The problem that critics have posed is to specify what these representation theorem results say about our actual beliefs or degrees of belief (Meacham and Weisberg 2011; Hájek 2008). What is this "*as if*" result?[20] The problem is a general one, not limited to only considerations of likelihood but also to the preference-based approach, where a utility and probability function is derived from an agent's preferences over acts or choices. Suppose that at an agent has preferences over acts such that there is a probability and utility representation of her preferences. What does this tell us about the agent itself? Meacham and Weisberg put it this way: "this gives us no reason to think that such agents *are* expected utility maximizers with probabilistic degrees of belief."(Meacham and Weisberg 2011, 5). They describe a situation where someone (they take Sherlock Holmes, but we need not be so dramatic) who is an expected utility maximizer with preferences that satisfy certain constraints, suddenly gets all twisted up mentally and has his utilities and degrees of belief all turned around, but whose preferences over acts remain the same. They say that this person (Sherlock Holmes struck on the head) can still be represented by a probability and utility function, even though his degrees of belief and utilities are all switched around; the representing probability and utility functions are too far detached from reality.

20. There is also the issue of whether the theory is to be construed as descriptive or normative; we have already assumed a normative interpretation here and will not discuss descriptive interpretations.

A similar point is also made by Hájek. "The concern is that for all we know, the mere *possibility* of representing you one way or another might have less force than we want; your acting *as if* the representation is true of you does not make it true of you."(italics in original Hájek 2008, 804) The complaint is that when the only requirement is that an agent's preferences over acts can be represented *as if* they maximize utility, that does not mean that the preferences over acts do maximize utility nor that the agent has utilities or probabilities. He makes the point by imagining representation theorem results for Voodooism: it might perfectly well be true that an agent, by his preferences, can be represented as if Voodoo spirits are warring inside him. But this hardly says anything about his actual mental goings-on. The challenge Hájek poses is to argue that it is probability and utility representation that matter and not voodoo-representations.

Drawing a distinction between different normative interpretations of representation theorem-results is helpful here. Different interpretations of these results differ in strength and can be classified as being anti-realistic or realistic in varying strength. These interpretations have been well investigated in Zynda 2000, and the interpretation of representation theorems based on considerations of likelihood are I will here defend is both weak-realism or antirealism—the difference is marginal, and of no consequence for our purposes.

1.3.1 Antirealism and Realism

Zynda distinguishes four different interpretations. The first is *eliminativism*, which is the view that (in our case) probabilities and utilities are unreal and simply do not exist. A position similar to this in philosophy of Mind with regard to all attitudes of folk psychology is eliminative materialism, as defended and argued for by Churchland (Churchland 2010). A slightly milder view Zynda calls *antirealism*, which is that the formal theory of probability and decision theory can or maybe should be adopted, but this is so generally on pragmatic grounds, *not* because the concepts used refer to existing things. There are no degrees of belief, probabilities or utilities, on this view; the concepts are useful in explanation, provide a cleaner theory, etc. If one says that degrees of belief or probabilities do exist one would be characterized as a realist. Realism is also subdivided into two positions: *weak realism* and *strong realism*. The weak realist says that degrees of belief (say) can be attributed to agents on the basis of actions, but that these degrees of belief or probabilities are only defined by reference to something more basic; say, elicited preferences over acts. In weak realism, probability exists, but only by logical construction of an agents preferences or considerations of likelihood. Strong realism, on the other hand, is the theory that degrees of belief, utilities and probabilities refer to genuine things. It says that agents have probabilities, have degrees of belief, and that coherent degrees of belief are probabilities.

There are good reasons for rejecting eliminativism and strong realism from the start—subjective probability does exist *in a way* and so do degrees of belief, but not in the strict way that we have probability functions in our head, surely. However, considering how these positions relate to the problems posed is still illuminating (and we should not reject any theory without first considering what would be the case were it true). For that, let us for the moment put Meacham and Weisberg's Holmes story to sleep and replace it with a less fantastical story of Zynda. Zynda (2000, 51–51) asks us to imagine to friends, Leonard and Maurice, of whom Leonard claims to have degrees of belief in line with the axioms of probability theory stated above, whereas Maurice claims to have something completely different which does *not* follow the probability axioms: degrees of belief that follow so-called *believability rankings* (Zynda 2000, 51–53). However, Maurice's actions do not violate what is required to be representable as an expected utility maximizer: Maurice uses a different method altogether for combining his degrees of belief[21] (which are *not* probabilities) with his utilities (which he does claim to have—the same as Leonard, in fact) to produce his preferences over acts; a measure Zynda calls *Valuation*. So, in the end, Maurice can be represented as maximizing expected utility and therefore, according to the theory, as having a probability function and an utility function. But he does not claim to *have* degrees of belief which are probabilities, whereas Leonard does claim to have these.

For the eliminativist there is no problem in explaining this thought experiment, for probability, degrees of belief, and utilities are all elements of a simply false and non-referring theory, a theory which ought to be discarded. Probabilities and utilities do not exist, nor are they part of a useful theory. Since we are in this thesis interested in obtaining a probability representation of degrees of belief, an eliminativist position is not an option. An antirealist position is to say that probabilities and utilities (and perhaps degrees of belief) do not exist but are part of a very useful theory. For the antirealist, because Maurice and Leonard in

21. Maurice claims to also have degrees of belief and Zynda has no quarrels attributing these (whatever they are) to him, but it is important that what are called "degrees of belief" differs fundamentally in Maurice and Leonard. As Zynda uses the term degrees of belief are psychological attitudes which are analogues of beliefs, but which might have all different sort of structures and follow different conditions—hence the possibility of believability rankings of Maurice as opposed to probability assignments to propositions. It is important in this thought experiment to distinguish "degrees of belief" and "probability". Zynda frequently talks about the realism of degrees of belief, but fails to clarify at points what degrees of belief are: only on a very strong conception of degree of belief are these numerical, and would they be probabilities.

the end have the same preferences, both will be represented as maximizing expected utility and therefore the same utilities (up to linear transformation) and probability functions represent them both, regardless of how the preferences came about. For the antirealist, what happens psychologically—what, in the thought experiment, the agents *claim* to do—is irrelevant and is described by Zynda as a "black box". The introspective testimony of both Leonard and Maurice amount to nothing; we need not and should not say that probabilities and utilities *exist* and we need not consider what degrees of belief *are*. What we look at for determining probability and utility functions is simply an agent's preferences and, if these preferences satisfy certain conditions, the agent can be represented as maximizing expected utility. Or, in the case of considerations of likelihood as we have discussed in this chapter, the antirealist is only interested in comparative judgements of likelihood, and through these can a probability representation be obtained. A choice between different models has to be made solely on the conventional or pragmatic virtues of the two rival models.

On both realist interpretations the story is quite different. For a strong realist interpretation, there is a fact of the matter about how Maurice came to his preferences, and a choice between a model based on expected utility and one based on believability rankings and so-called valuation is to be made, quite simply, on what exists; what psychological structure is the true one. So, if the introspective testimony of both our characters is correct, Leonard has probabilities and utilities, and maximizes expected utility, whereas Maurice does not. Of course, the strong realist is not tied to introspection, and Maurice could just as well be said to have probabilities and utilities; the essential point is that for the strong realist both these notions refer to genuine psychological attitudes or things and are not abstract things defined by construction from, say, one's preferences. Agents *have* probabilities, utilities, and degrees of belief.

The difference with weak realism is that for the weak realist probabilities and utilities are abstracta. For the weak realist, probabilities and utilities can be ascribed to an agent and they do exist, but they are not things or attitudes that exist without considering the theory they are a part of. So probabilities and utilities are said to be real, but the notions are defined by logical construction from preference rankings over acts (or something else, like considerations of likelihood). As Zynda notes, this is analogous to antirealism with respect to, say, centers of gravity (Zynda 2000, 60). A center of gravity is an abstract concept that can be said to really exist, although it makes no sense to talk about centers of gravity outside of the framework of physics. A center of gravity is real insofar as all the gravitational forces that attach to some object are real, but it is the point where the resultant torque of these forces is zero; without taking these forces acting upon some object into account the notion of center of gravity becomes meaningless. On a weak-realist position, notions are logically constructed from other elements—in our case, preferences or considerations of likelihood—and probabilities and utilities are then defined to be real, but only by this construction.

For both forms of realism there is a fact of the matter of which of several competing theories and descriptions is the true one; a choice between competing theories is to be made on considerations of the truth of the theories. This is not to say that both versions of realism cannot resort to pragmatical virtues of theories such as simplicity, explanatory value, etc; realists can argue for the existence of attitudes and truth of theories based on some theory being the best explanation, and frequently do. But their argumentation is mostly aimed at proving the existence of something and the truth of a theory; the antirealist does not take this extra step towards some theory being the true theory.

Weak realism differs here from strong realism in that its definition of probabilities and utilities (and possibly degrees of belief) by construction from primary terms makes the argumentation with regard to Zynda's thought experiment slightly easier. The strong realist wants to argue (to be convincing) that Maurice's theory is false or weaker than the standard theory based on expected utility. Weak realism on the other hand has an extra path open to her, which is to say that probabilities and utilities are defined by things that are equivalent in both Leonard and Maurice—despite Maurice's best introspective intentions—and thus that Leonard and Maurice in fact have the same probabilities and utilities.

This is done by saying that what is real are the considerations of likelihood and that this is all there is to degree of belief; if Leonard and Maurice both share the same considerations of likelihood, and these considerations both satisfy the conditions discussed earlier, they can be represented by the same probability function. Thus, it is to define a degree of belief as a qualitative consideration of likelihood. In the case of expected utility theory, it is to say that what exists are preferences over acts, and that probabilities and utilities exist only by construction from these. The self ascriptions of Maurice and Leonard therefore amount to nonsense: probability and utility are representations that are based on preferences, and the preferences over acts of both Maurice and Leonard are the same.

This weak-realist position is the position Zynda also ascribes to. His reason for choosing this position over antirealism is that the concepts of folk-psychology are so widespread in usage that not accepting the things it talks about is implausible (Zynda 2000, 66). But, in our approach to qualitative probability, it should be noted that we do not postulate *degrees of* beliefs, as in numerically-valued beliefs, nor probabilities and utilities, for these are only mathematical representations. So this type of realism is minimal: it only postulates the reality of beliefs insofar as these are required to obtain a probability representation, or preferences insofar as these can be represented by a probability and utility function.[22]

Barring possible realistic inclinations—a taste for realist theories over antirealist ones—this weak-realism does not differ very strongly from an antirealist position. For all the weak-realist posits are those things that the antirealist might not have any problems with: preferences over acts and/or considerations of likelihood of events. The only difference, then, is that the realist says that the probabilities and utilities constructed from these exist, where the antirealist says they do not; the difference is minor and of no consequence for our purposes. It is not clear from this construction alone what it means to say that probabilities and utilities exist or not; they are constructed from something unproblematic and therefore the difference between their existence or nonexistence arises from other considerations. For example, semantic considerations of the meaning of truth. The only question relevant here is what the precise definition of degrees of belief as considerations of likelihood or preferences, is—must they be measured values, dispositions to elicit values? This is a question we will return to in chapter 3.

1.3.2 Is this weak-realism strong enough?

Despite the little difference between a weak-realist and antirealist account, let us for now call our position "weak-realism" which ignores the difference between weak realism and antirealism. The realistic portion of the construal are the considerations of likelihood; probabilities are defined by these. Whether we call probabilities "real" and say that they exist has no consequences. So on this interpretation representation theorems do not establish that we *have* precise degrees of belief and probabilities (and utilities), in the sense that these are genuine mental attitudes, but only insofar as we have considerations of likelihood,

22. Zynda says to be a realist with regard to comparative judgements of likelihood of propositions, but he makes remarks to the extent of being an *anti*realist with regard to preferences (Zynda 2000, 66–67). I cannot understand this distinction: preferences over acts are also understood as an ordinal structure, where acts are compared to one another on the basis of preferences. To not allow these comparative preferences while allowing comparative judgements of likelihood seems ad hoc. Other remarks of Zynda on the same pages also seem to contradict this position; I therefore ignore this distinction and allow realism for preferences, too.

through which probabilities are defined and determined. If we do, this is not in virtue of representation theorems. We want as little structure on beliefs as necessary to obtain a probability representation to make the approach as plausible as possible—precisely because beliefs are not straightforwardly numerical, as argued above and in chapter 3. Let us now return to the problems posed by Hájek and Meacham and Weisberg.

On the weak realism explained, the possibility that an agent would have her utilities and probabilities twisted up but with the same preferences, or have her probabilities all changed with the same considerations of likelihood, is conceptually impossible. Probabilities and utilities are representations of rational or coherent preferences, so if the preferences remain the same, so do the functions. The same goes for a part of Hájek's criticism: agent's *do not have* probabilities, but are represented by probability functions.

Meacham and Weisberg are not satisfied with this view (Meacham and Weisberg 2011, 21–23). First, they object that being a realist about an ordinal structure is almost just as implausible about being a realist about degrees of belief. This claim seems unjustified: an ordinal structure is much more plausible, for it says only that agents compare elements of a certain set of propositions in likelihood. That is surely less of a theoretical postulation than talking about genuine degrees in one's head. Secondly, they desire an interpretation of representation theorems which *does* say that agents have degrees of belief, possibly vague; at the very least, it must be a stronger structure than just that of a qualitative ordering, for they want also to say *how much* more likely some proposition is than another. Frequently, they say, we compare our own views on matter on more than just an ordinal scale, and thus that our beliefs have a stronger structure than just a qualitative ordering of likelihood. There *is* a magnitude in our *beliefs*, not just in our representations. Or so they maintain.

I think this criticism is unfounded. First, all the weak-realist interpretation of the representation theorem says is that the sufficient conditions for degrees of belief to be representable is to have a certain structure and to satisfy certain conditions; our beliefs might be stronger structured than is necessary for the representation, but that is not *required* for a probability (or utility) representation. One might be a realist about numerical degrees of belief for other reasons than the representation theorem result. Secondly, the stronger the structure present in our beliefs and the stronger the structure we posit among beliefs, the easier it is to get a representation of it: if the structure of qualitative likelihood is Archimedean or can be reduced to such a structure (as is the case in (Suppes and Zanotti 1976)), a representation is obtained with much less conditions. The more structured out beliefs, the weaker the axiomatic system need be (Krantz et al. 1971, 208). Scott's axiom also imposes strong structure on beliefs. If we adhere to a weaker structure, a stronger axiomatic system is required. Now if we *would* have numerical degrees of belief, then a representation theorem is quite easily established as we shall see in the chapters to come: all that is required is boundedness and additivity. This does not make probabilism anything like a trifle; the justification for the conditions is not easily established. Lastly, Meacham and Weisberg's claim that our use of language and belief-ascriptions seems to indicate numerical beliefs to some extent is ill-defended; this usage is surely much too vague to be taken to indicate our having genuine degrees of beliefs in our head. Every such ascription of degrees is, I maintain, very rough and just guessing; regardless of whether or not we say that we believe something to a degree of 0.3.

Now we can also return to their Holmes-struck-on-the-head story. Let us replace the preferences over acts with considerations of likelihood over events to better fit the rest of this chapter; it does not matter for their thought-experiment. Just before Holmes was struck on the head, he had a certain ordering of likelihood of events, which we used to calculate and determine a probability function that represents this ordering. Then someone struck him on the head hard which, according to Meacham and Weisberg, messed up his probabilities even though his considerations of likelihood remained the same. But the thought experiment is impossible. If probabilities are determined by considerations of likelihood, the former cannot differ if the latter do not. Probabilities and utilities cannot float free from preferences, either. Nor is it even meaningful to assume that coherent considerations of likelihood can float free from degrees of belief, as their thought experiment can be understood to presuppose: even if there were precise degrees of belief, surely a high degree of belief in a proposition would mean considering that proposition to be very likely. Similarly, preferences over acts in situations of uncertainty also involve considerations of likelihood of states of affairs; having a high degree of belief in some event being the case, and attaching a high value to some outcome, implies preferring the act that is a function from that event to that outcome, to an act that is based on a less likely event with a less valued outcome. To not have such preferences is irrational and incoherent—in which case there is no utility or probability representation.

As for the problem of additional representations such as Voodooism: on the interpretation expounded we do not get the result that being representable in a specific way allows us to infer that we *are* that way. And if Voodooism is posed as an alternative to probability theory and utility theory then, quite frankly, it will not take much investigation to see which of the two theories is more explanatorily successful.

To sum up, representation theorem-based arguments then try to get a representation of degrees of beliefs, and what is required for this is completely dependent on what structural assumptions we make on these degrees of beliefs. The problems with the approaches so far is that the normativity has not been really argued for, and that the conditions are especially implausible for beliefs purely as such—the interpretation that seems most natural for epistemic probabilism. If we talk about *measurement*, stronger structure of degrees of belief seem more plausible, making a representation theorem easier to argue for. And the theories famously created by De Finetti and Ramsey talk exactly of degrees of belief as *measurable*; it is to these we must now turn.

Chapter 2

De Finetti's construction of subjective probability

The approach of this chapter is to obtain a representation of a structure stronger than the noncommittal considerations of likelihood we used in the previous chapter. The approach can be described as starting from dispositions to act in situations of uncertainty and trying to obtain a probability representation of these dispositions to act. The approach proceeds from the idea that if we try to obtain a probability representation of degrees of belief, these degrees of belief should be *measurable*. And the way of measuring beliefs is through actions: if I am pretty confident in something being the case—I consider it likely—I am disposed to act in ways different from when I would be less confident. Examples are abundant: I consider it more probable that a Cola can will not spray than that it will, provided it has been treated normally, and am therefore not too careful in opening them. Considering something likely makes it that I am more disposed to choose an act based on that likely event rather than on an event I do not consider likely.

The work of Ramsey (Ramsey 1926) and de Finetti (de Finetti 1993, 1980, 1974) has been very influential and provides a very clear and systemic approach to a probability representation of degrees of belief. Utility is taken to be fixed and probability is defined over dispositions to act in certain particular measurement scenarios. In one of these scenarios the degree of belief in some proposition is defined as the price for which an agent would be willing to both buy and sell betting contracts conditional on that proposition being true or false, and in another the degree of belief is defined as a forecast of a random variable, given that this forecast will be penalized according to its distance from the truth. In the preceding example of me opening a coke can, this means that I would be willing to bet a fairly high price that the can will not spray all over my shirt, compared to that it will, or forecast it pretty likely that the can will not spray. The general idea behind this is that you would stake more on events you consider likely then on unlikely events. This scenario also gives coherence conditions on these prices, where all and only coherent prices can be represented by means of a numerical probability function. Let us now look at de Finetti's[1] setup and definition of probability, which we will discuss and evaluate with respect to epistemic puritanism in the next chapters.

2.1 De Finetti's setup

De Finetti (1974) formulated his theory in the general case of random variables, which we have already introduced. We will from now on restrict ourselves to a finite sample space Ω due to the technical difficulties related to countably infinite partitions. This restriction is artificial but does not alter any of the points made in this thesis.[2]

We start, as before, with a partition Ω of possible 'points' or 'atomic' events, of which only one point or value is the true one for an agent. This is also the space over which the random quantities are defined. Events are subsets of Ω; more specifically, the set of events $\mathcal{A}$ is an algebra over Ω. De Finetti emphasizes here that this set of possible points is *not* to be taken as being fixed or as having a specific structure: all that is required is that it is some partition of events of which one and only one is true. The usage of the term "points" is not to be taken to mean that Ω is some fundamental set which cannot be subdivided. How this sample space is divided depends on any situation of uncertainty for an agent and what possibilities the agent considers in the decision problem (de Finetti 1974, 33–34; 1975, 267–276).

2.1.1 The space of possibilities

However, something de Finetti does assume about any situation of uncertainty, is that the true point can be verified: every outcome for a well-defined decision problem must be *verifiable*. This means that in any situation of uncertainty for an agent what is in fact

1. Due to Ramsey's death at 26 years old, de Finetti's version of the theory is the much more developed one. For this reason, de Finetti's setup is the one used in this thesis, although we will also discuss aspects of Ramsey's account in the next chapter.

2. See (Williamson 1999; Howson 2008) for two important papers on this topic.

the case (what is the outcome) must be verifiable. There must be a fact of the matter about every event or proposition and it must be well-determined along with us being able to determine what this fact of the matter is. This is not to say that something must be *immediately* verifiable, but only that we could, possibly, verify it.[3] This condition is partly imposed by de Finetti's operationalism—in the measurement scenarios discussed below we must be able to penalize elicited 'probabilities', and every betting contract must be decidable (de Finetti 1974, 28)—but de Finetti also defends this from his conception of meaning: only if there is a method of verifying does it *make sense* to talk of an event (de Finetti 1975, 264–267). How exactly this notion of verifiability must be cashed out is something we will not discuss here.

With this space of *possibility* in place, we can introduce probability and prevision. De Finetti starts by introducing a particular kind of random variable: the random *gain*. These are also sometimes called "gambles": the gain is dependent on what point w of Ω is the true point and can be understood as functions from Ω to $\mathbb{R}$, where the numbers of $\mathbb{R}$ here represent units in some measure of linear utility. The simplest such measures of linear utility are monetary values. Money, of course, is not valued linearly across the board by most people, where the difference of value between a million euros and a million and one euros is negligible. In small quantities of money in relation to the agent's wealth, though not trivial, it seems reasonable enough to hold that money can be said to be of linear utility for the agent.[4] An agent then has an expectation of this random gain, which can be defined as the *price* she would be willing to exchange for this random gain. This interplay of random and certain gains can be understood as follows: an agent considers the random variable and has an amount of money in her pocket. We can imagine asking this agent how much she is willing to pay for this random gain: is €0.05 ok? How about €50? For low prices the agent is willing to exchange the certain gain for the random gain,

3. I take this also to be de Finetti's position, because he writes "[we think] of the fact of whether or not we can obtain the information that for us determines whether it is true or not (or, at least, whether there is a *possibility*."(my emphasis de Finetti 1975, 265) Peter Walley, however, thinks that de Finetti identifies 'meaningful' with 'observable' (Walley 1991, 466); I think this interpretation is wrong. In any case, the weak notion of verifiability expounded here is, I think, still strong enough for the setup of de Finetti.

4. Walley constructs a linear utility scale by using a lottery with a large number of tickets, all with an equal chance of winning a certain prize (Walley 1991, 59). This has the advantage that the number of lottery tickets is (plausibly) of linear utility: if the number of tickets an agent can gain on the basis of some outcome is greater than some other, the utility of that outcome is also greater by that proportion. The problem with this setup is that it defines probability by using a lottery of which every ticket has equal probability.

but she is not so confident in the random gain to be willing to exchange €50 for it. This is because she is sure enough that the actual value of the random gain will turn out high enough to exchange the small price, but not sure enough to exchange big amounts of money. Assuming a complete preference relation between random and certain gains, any random gain higher than some certain gain is desirable for the agent, but any random gain lower than some certain gain is undesirable. For every random gain, there is then some highest certain gain for which the agent is indifferent to the exchange: this is the *fair price* $\mathbf{P}$ of X for the agent.[5]

This fair price $\mathbf{P}$ is a *representation* of the "opinion of an individual who is faces with a situation of uncertainty"(de Finetti 1974, 83). The fair price is a mathematical function which has the following properties (de Finetti 1974, 74):

- Additivity: $\mathbf{P}(X + Y) = \mathbf{P}(X) + \mathbf{P}(Y)$.
- Boundedness: $\inf X \leq \mathbf{P}(X) \leq \sup X$.

For the special case of probability, this comes down to that, in the case of two incompatible events A and B, $P(A \vee B) = P(A) + P(B)$, plus that for any A, $P(A)$ should never be lower than 0 or higher than 1. Also, any partition must always be given a total probability of 1; that some element of the partition happens is certain, and the total partition must therefore be given probability of one. These fair prices $\mathbf{P}$ are the same functions that are allowed by the axioms of the probability calculus, with the only difference of de Finetti's assumption of finite additivity. Thus, probability is *defined* as the fair price of an agent for some random variable. $\mathbf{P}$ is defined as a representation of an agent's preferences over exchanges among certain and random gains, where the properties of $\mathbf{P}$ correspond to assumptions about these exchanges. Fundamental is the hypothesis or assumption of *rigidity in the case of risk*, meaning that de Finetti assumes that the agent's preferences can be represented as linear.

So far, the theory is a theoretical representation of preferences, rooted in a decision theoretical picture. This theory by itself does not have strict empirical significance, and we are unable to say when an agent does and when he does not follow the conditions of the price function in his preferences. The fair price $\mathbf{P}$ is intended to represent "the opinion of an individual who is faced with a situation of uncertainty" (de Finetti 1974, 83), but so far de Finetti considers the theory only "metaphysical-verbalistic": he requires a definition

5. Although the fair price is always of some random variable and should properly be written as $\mathbf{P}(X)$ (for any X), we will disregard the random variable in the notation and simply write $\mathbf{P}$ for the fair price.

which allows for *measurement* of this probability (de Finetti 1974, 76). He accomplishes this through the construction of two measurement scenarios. These measurement scenarios allow for an *operational* definition of subjective probability through an agent's acts; it gives empirical significance and empirical meaning to the notion of probability. The properties of price (or probability) functions **P** define what coherence actions are in the measurement scenarios. Moreover, the scenarios de Finetti proposes allow for the understanding of the coherence conditions on elicited values as rationality conditions. Incoherence, in these scenarios, amounts to being open to the possibility of sure loss: incoherent elicited values suffer an unnecessary high penalty; coherent values do not.

2.1.2 The betting scenario

Let us start with the betting scenario. The random gains in this scenario are betting contracts which yield a gain or loss dependent on the actual value of the random variable. The (fair) price of an agent is the price at which the agent is said to be willing to both buy and sell betting contracts. The gain of these betting contracts is the actual value of X minus the elicited value $\bar{x}$; this gain can also be negative, so that the opponent in the scenario—the one who buys or sells betting contracts from the agent—gains this amount. This is done by having the opponent determine the stake c; if c is negative, the gain of the agent is $\bar{x} - X$.

The elicited values $\bar{x}$ are prices for betting contracts, but they cannot always be represented as a price function **P**. This is so only if a set of prices elicited by the agent is coherent: a set of prices is coherent if there is no finite series of betting contracts bought and sold at elicited prices $\bar{x}_i$ for which the agent would lose no matter what value X_i takes, for all $i = 1, ..., n$. Thus:

$$Y = c_1(X_1 - \bar{x}_1) + c_2(X_2 - \bar{x}_2) + ... + c_n(X_n - \bar{x}_n)$$

In the case of incoherence, there is a set of stakes $c_1, ..., c_n$ for which the random gain Y is always negative; if coherent, there is no such set of stakes. A (finite) series of betting contracts that would lead to sure loss is called a Dutch book. As it turns out from de Finetti's theorem of total probability (de Finetti 1974, 99–101) which I will not here prove, these coherence conditions of elicited betting prices are exactly what is needed to ensure a probability representation: the coherence conditions on elicited prices correspond to the properties of **P** listed above.

These coherence conditions on elicited values can also be expressed as the condition of being *undominated*, which is another clear way of presenting the (in)coherence conditions.[6] Remember that the agent is asked to state her price for gambles X_i, defined over the algebra $\mathcal{A}$ over Ω. These prices are, in effect, *acts* in the sense of Savage (1954), and dominance is as such defined over *outcomes* of posted prices; these outcomes are gains and losses. In this criterion of dominance, posting a set of prices is compared to the act of *abstaining* from entering any bets. Abstaining here is defined as the act of which the outcome is 0: the agent neither gains nor loses money in the act of abstaining.

A set of announced fair prices is then said to be dominated if there is a series of betting contracts based on these prices for which the *net* payoff is lower than abstaining, no matter what happens. And if a set of fair prices is dominated, it is incoherent. If for every set of betting contracts based on the fair prices there is some outcome of the random variables for which the agent does better than abstaining, the set of fair prices is undominated and hence coherent.

2.1.3 Scoring Rules

Scoring rules were introduced by de Finetti later on and preferred by him over the betting scenario, because of the absence of an opponent in the scenario.[7] The idea is similar, but there are essential differences. Essentially, scoring rules are methods of evaluating accuracy of expectations or *forecasts*, where a given expectation of any random variable is given a penalty proportional to the distance from the actual value of that random variable. Thus, in the case of an event, a forecast is a number that is penalized according to the distance to the actual value of the indicator function of that event: 0 if the event does not happen, 1 if the event happens. The general idea as an elicitation procedure is to ask any agent what her forecast of any random variable is, given that she will be penalized according to the square of the difference from her forecast of the random variable and the actual value the random variable takes.

The specific kind of scoring rule de Finetti used here are strictly proper scoring rules. Scoring rules are random quantities f, g which are functions from expectations or forecasts and random variables X, where the value the scoring rule takes represents a penalty in some measure of utility. Thus, if a certain set of forecasts or expectations gives a higher score than another set, the *latter* set is preferable, for the scores are penalties and it is

6. In this and the discussion of scoring rules I draw from a presentation of Seidenveld (Seidenveld 2014).
7. See especially (de Finetti 2008, 1–30; 1981).

desirable to minimize the penalty. It is important to highlight here the difference between a forecast and expectations, because the actual expectation of an agent of X need not be the elicited forecast; the agent might name some other value then she actually expects. However, strictly proper scoring rules make it in the agent's interest to elicit her actual best expectation, and we can therefore use only the notation $P(X)$ for the forecast of an agent of X.[8] The strictly proper scoring rule de Finetti uses is the Brier score: $L = (X - \bar{x})^2$.

Strictly proper scoring rules also make it in the agent's best interest to be coherent in the forecasts, where the coherence conditions are the same as in the betting scenario. It works as an elicitation scenario as follows: an agent is asked what her forecast or expectation of some random variable is, being told that she will be penalized according to the square distance from her announced forecast and the actual value of the random variable. This penalty is, again, in some measure of utility. If the agent is incoherent in her announced forecasts, she opens herself up to an unnecessary high penalty: there is a set of forecasts for which the score is uniformly less. So whatever happens, whatever value the random variable takes, she could do better in some way—could ensure a lower penalty. By contrast, coherent forecasts do not have this property: there is no set of forecasts which are uniformly better than any set of coherent forecasts. So incoherent forecasts are *dominated* by some other set of forecasts, whereas coherent forecasts are undominated. Dominance means that there is a set of forecasts for which the net score is lower no matter what happens.

There are important differences between the scoring rules elicitation scenario and the betting scenario: for one, in the latter we had prices for gambles, whereas with scoring rules we have forecasts which are not (necessarily) monetary or in some other measure of linear utility. Another important difference is that in the betting scenario an opponent is required to elicit prices and to buy and sell gambles; this is not required for scoring rules. Elicitation by means of scoring rules simply requires a forecast which is penalized; no opponent is present in the penalization and the penalization is strictly based on the elicited value. However, the elicited value does depend on the penalty: it is because the scoring is strictly proper that there is a single elicited value which is the agent's best estimation of a random variable. Also, dominance in the betting scenario is characterized by comparison

8. A *proper* scoring rule is a scoring rule where the prevision of the penalty of $P(X)$ is at most as high as the penalty of some elicited value. More formally: $P(f(P(X_1),...,P(X_n))) \leq P(f(\bar{x}_1,...,\bar{x}_n))$. This means that the *expected* penalty is lowest for the best expectation of X of the agent. A scoring rule is *strictly* proper if $P(f(P(X_1),...,P(X_n))) = P(f(\bar{x}_1,...,\bar{x}_n))$ only if $P(X_i) = \bar{x}_i$ for all $i = 1,...,n$. This says that it is in the agent's best interest—assuming she wants to minimize expected loss—to elicit her actual expectation of a random variable. Any other value will have a higher expected penalty.

with abstaining, where the loss could be uniformly less if the agent would abstain from betting, compared to some set of betting prices. For the scoring rule scenario, there is no abstaining of such a sort: dominance is simply that some set of forecasts is possible which yield a lower penalty (lower score). The important thing, though, as regards dominance is that the coherence conditions on forecasts and prices in both scenarios are the same: both correspond to the properties of $\mathbf{P}$. Thus, if and only if one's forecasts or prices are coherent is there a probability representation of these values.

2.2 Operationalizability

We have already mentioned de Finetti's requirement of operationalizability, where the two scenarios are intended as genuine measurement scenarios of an agent's personal probability. The setup of the price of an agent is a setup of the *operational* definition of probability as consistent betting prices, and the scoring rules define also probability as an elicited forecast. In both scenarios the probability is assumed (by the way the measurement is constructed) to be a precise value. Whether this is justified we will discuss briefly in the next chapter; for now let us focus on the role of this measurement and the complications it involves.

Even though probability is defined as coherent betting prices (in the betting scenario) or coherent forecasts (with respect to scoring rules), this is not to say that any intuitions we have about "expectations", "degrees of belief" or "probability" are irrelevant; nor did de Finetti state any such thing. It is important to ask if *what* is being measured in a measurement scenario is actually an agent's degrees of belief or best expectations—putting aside what these degrees of belief are until the next chapter, let us here use the term "degree of belief" as meaning an agent's actual opinions about events. We will not assume these to be numerical but take these to be only considerations of likelihood: if my probability for some event, as defined in the setup, is high, so should my degree of belief in the respective proposition be high.

The main problem in the betting scenario is the presence of an opponent in the scenario. Suppose that someone whom I know to be no expert on road cycling, wants to measure my degree of belief that Tony Martin wins a certain time trial. (For those readers that are not road cycling enthusiasts: Tony Martin is a multiple time-trial world champion on road cycling.) I actually know that this opponent is quite misled by a friend, such that he considers the likelihood around 0.3; I consider it to be around 0.7. *Knowing that his* *considerations are ill-informed*, I post a price of 0.5, which he thinks is a good deal: he thinks the odds are lower, so he will gladly sell the betting contract so that, as he thinks, he can pocket €0.2 (ignoring the stakes for convenience, although these prices are a triviality). However, the advantage is mine here, for I bought a betting contract for a lower price than my actual price. So it is in this situation *in my interest* to post prices that are *not* my actual opinions and not my actual degrees of belief, and I am also coherent in posting these different prices.

Similarly, (sports providing such easy examples,) if I were asked a year ago for the probability that a low ranked American Football team would win against the Seattle Seahawks I, as a complete novice on the subject, might as well have posted 0.5. Suppose that my opponent, however, is an expert and *as soon as I post 0.5* he happily sells me that bet: he takes the probability to be about 0.1. In noticing his reaction, I become aware that my price was ridiculous and that I just lost money by overpricing the bet; I change my degree of belief immediately. So *just being measured* can change my degree of belief; this is an undesirable result.

These aspects are absent in the forecasting scenario which uses scoring rules; here there is no opponent who can influence the agent's degrees of belief. It is only the agent whose forecast is being measured; there is no opponent whose state of information can influence this forecast. De Finetti, as an operationalist, defined subjective probability as values elicited in one of the measurement scenarios. Because of the problems with the betting scenario just discussed, he preferred the definition of subjective probability by means of scoring rules. However, there are other problems with a purely operational definition. The most important one is the generality of any operational definition: why should an operationally defined measured value be also a *general* degree of belief or a general theory of decisions? Such a general theory is desirable from the purely epistemic standpoint, and an operational definition is inadequate for the epistemic puritanist. They want a general theory of subjective probability theory as about general degrees of belief, not just a theory of the specific cases where these beliefs lead to measured values.

Yet we can extend the interpretation by going hypothetical: it is not about *actual* acts and measurements, but about *possible* acts. More precisely, it is about dispositions to act in measurement scenarios; either in an ideal betting scenario or in the scoring rule scenario, for both yield the same admissible degrees of belief (the coherent ones). This more theoretical interpretation is to say that the prices of agents are models of how an agent *would* bet in a hypothetical circumstance. This also allows for the setup that defines probability as price functions to avoid the problems of *actual* measurement and opponents just described, by saying that probability is what a rational agent would bet in ideal and hypothetical betting scenarios. This interpretation seems to differ from the theory of de Finetti and is more like that of Ramsey, who also emphasized the definition as dispositions to bet (Ramsey 1926, 34). However, as an actual measurement scenario scoring rules have preference over the betting scenario, and degrees of belief can also be defined as dispositions to forecast when subject to a scoring rule.

That dispositions to bet or accept betting contracts is actually a quite general characterization of dispositions to *act* can be seen by considering Ramsey's statement that "all our lives we are in a sense betting" (Ramsey 1926, 42). The difference between that of the betting scenario and the general case is that the 'bets' (that the subway comes on time, that this Cola-can won't spray all over my shirt) do not yield payouts in monetary values but in other measures of utility. I value taking it easy in the morning considerably, and therefore take the chance that the subway might not come on time (in which case I would be late for class) to be worth it. Monetary values are only a specific and linear form of values, represented by utilities. The idea, therefore, is that the agent is *disposed* to bet, and that these dispositions are, if coherent, an agent's subjective probabilities. What these dispositions are we will discuss in the next chapter.

In this chapter we have discussed de Finetti's setup and construction of subjective probability and discussed only some problems regarding these, and have so far not discussed whether these probabilities (and price functions $\mathbf{P}$) are degrees of belief and whether they can be adapted for a purely epistemic Bayesian Epistemology. It is to this topic that we must now turn.

Chapter 3

Prices and Degrees of Belief

De Finetti's theory of subjective probability has been both heavily criticised and acclaimed. The benefits of it are that it successfully creates a theory of subjective probability: probability functions (or, more generally, prevision functions) are defined as representations of coherent prices for random gains, representations of elicited values in one of the measurement scenarios. The idea of the betting scenario is surprisingly simple, boiled down to its essence: how much are you willing to bet on this event? On top of that, it also fleshes out the normativity of the coherence conditions: it is in an agent's best interest to be coherent, to minimize expected loss by avoiding sure loss.

But it's not all a success story. In the previous chapter we have seen price functions **P** and taken these to be representations of elicited values in measurement scenarios, but what these prices for betting contracts—in the case of the betting scenario—are supposed to be is not exactly clear. On a strong operational position, prices are representations of *only* elicited values in measurement scenarios. This might be of some interest, particularly in evaluating the opinions of experts[1], but a more general and less strongly operational interpretation is desirable. We have also established that a weaker and more general interpretation is taking prices as *dispositions to act*; still, the worry is whether this is sufficiently general and, more importantly, whether these dispositions to act can be taken as purely epistemic *degrees of belief.* In line with the idea of chapter 1, many epistemologists have interpreted de Finetti's setup as being about degrees of belief and have the **P** functions

1. See de Finetti's discussion of scoring rules as applied to the exploration of oil deposits, where the opinions of experts are made precise by asking them to forecast in a scoring rule scenario (de Finetti 1972, 4–6).

be representations of coherent degrees of belief. This epistemological interpretation is the foundation of what we have called Bayesian Epistemology: the theory says that beliefs come in degrees and it cashes out coherence in an equivalent fashion as de Finetti's setup: it is irrational in some way. The purely epistemic interpretation then interprets the theory as being a purely epistemic theory, where values should not be part of the picture and setup. However, the problem is that agents' preferences and values play a major role in the setup, both in the definition of degrees of belief and subjective probabilities, and in the measurement scenarios and the associated coherence (rationality) conditions. For this reason, others have argued that the setup is epistemologically vacuous: monetary penalty and betting prices have nothing to do with epistemology.

In this chapter we investigate what fair prices represent; in particular, whether they can be taken to represent degrees of belief and used as the foundation for Epistemic Puritanism Bayesian epistemology. This second question hinges on many (interconnected) properties of beliefs: whether beliefs can plausibly be taken to have enough structure for the coherence conditions to apply, what role the values play in the setup and whether they can adapted or removed, and whether rationality is or can be made an epistemic defect. As shall become clear, such an epistemological interpretation is very problematic, for in de Finetti's setup, price functions and previsions represent acts, not beliefs—at most, dispositions to act. This connection with preferences and values is undesirable from the standpoint of Epistemic Puritanism. In this chapter I argue that subjective probability and degrees of belief or prices are, in de Finetti's setup, essentially connected to values and preferences and that the purely epistemic standpoint has a very hard time defining or characterizing degrees of belief without these values and preferences.

3.1 What are prices or forecasts?

Let us start with considering some features of the price functions **P**. The most important feature of prices in the last chapter was that they are some precise numerical value. This value is, as we have seen, a monetary value or some other measure of linear utility: it is called by De Finetti a "certain gain", for which an agent is indifferent to an exchange with some random gain. We have seen that the fair price is constructed by a comparison of random gains and certain gains, for any individual; it is "inserting the degree of preferability of a random gain into the scale of certain gains", which De Finetti calls a "prerequisite condition of all decision-making criteria" (de Finetti 1974, 73). De Finetti called the function **P** a *price* and not a probability because it is the highest price for which an exchange of a bet with an amount of money is still acceptable to the agent. This price is exactly what the betting scenario is all about; talk of prices in the measurement scenario of a scoring rule is misplaced, where there are *forecasts* which are (represented by) probabilities, that are not values in some measure of utility. (The *penalty* of the forecasting scenario *is* in some measure of utility.) Forecasts and betting prices are, however, equivalent in that the same values are characterized as being coherent. It is worth focussing on the price in detail, though, because of its foundation in a decision-theoretic picture.

The price of an agent is called a "fair price" because it is the price for buying *and* selling: the agent would be willing to trade either way. The additive property of **P** makes it so that a (finite) *sum* of random gains is then fair with respect to a function **P**, only if each individual gamble is fair with respect to this function. This is De Finetti's hypothesis of *rigidity*: additivity of the price of a random gain of any agent. As a property of **P**, it is only a mathematical property, but as a property of an agent's willingness to exchange money for bets, it is a property of an agent's values. This rigidity of betting prices is hardly ever really the case across the board—I might not be willing to pay €1000 for an event A just because I would be willing to bet €100 for every one of the ten independent events that constitute this event A, simply because €1000 is too much money for me to bet on pretty much anything. At some amount of money the element of risk might influence the price for a bet (along with other considerations). Rigidity is applicable in practice only when the stakes of bets are chosen appropriately (small in relation to the agent's wealth, but not trivial).

3.1.1 Degrees of belief and measurement

Because of Ramsey's formulation of the theory as about degrees of belief, explicitely, Ramsey's (1926) construction of subjective probability theory is important and informative here. His setup is largely identical to De Finetti's, with the essential difference of being constructed around the notion of utility instead of monetary values. Ramsey stipulates in the beginning of his construal that whatever the notion of degree of belief might turn out to be, it has to be measurable. Only if measurable can we meaningfully ascribe degrees to beliefs to agents. After this, he also quickly rejects that these numbers can be assigned by means of introspection—this would not be a precise way of assigning numbers. The correct way is in measuring degrees of belief through behaviour; belief as a basis for action (Ramsey 1926, 34).

This might sound a lot like a behaviourist interpretation, which is a sort of interpretation no longer in favour in philosophy; the theory has been largely abandoned and is seen as an untenable position. Just a few of the problems related to a behaviouristic conception of "belief" are that many beliefs never lead to behaviour, many different beliefs can cause the same behaviour, and that explanation of behaviour is usually done by referring to beliefs, leaving a conceptual reduction impossible. To illustrate this, let us consider the question that Davidson asks: what behaviour corresponds to the belief that there is life on Mars? If this is said to be an utterance, "this shows he believes there is life on Mars only if he understands English, his production of the sound was intentional, and was a response to the sounds as meaning something in English; and so on. [...] [No] matter how we patch and fit the non-mental conditions, we always find the need for an additional condition (provides he *notices, understands*, etc.) that is mental in character."(Davidson 1980b, 217) Such a definitional behaviourism is circular and always requires the concepts it is attempting to define in its definitions.

Ramsey briefly discusses some of these problems. He first discusses the problem that many beliefs do not lead to action; many beliefs cannot be observed, let alone measured. Ramsey parries this worry by noting that it is about a *disposition* to act: it is about beliefs that "*would* lead to action in suitable circumstances; just as a lump of arsenic is called poisonous not because it actually has killed or will kill anyone, but because it would kill anyone if he at it."(Ramsey 1926, 33, emphasis added). So according to Ramsey, for degrees of belief it is required that every degree of belief can, in suitably ideal (hypothetical) circumstances, lead to action. As we shall see later, the measurement scenarios are just such circumstances. The second objection Ramsey discusses is that many different beliefs can lead to the same effects even in these suitable or ideal circumstances. Ramsey's short reply here is that it is the effects we are interested in, not the mental causes. So even though Ramsey grants that it is true that many beliefs can lead to the same actions, we are only interested in the actions: differences in beliefs are unimportant in these measured scenarios.

In saying this, Ramsey in effect *defines* degrees of belief as possible actions: degrees of belief are dispositions to act. Different (degrees of) beliefs that would lead to the actions in ideal situations are treated as identical. Ramsey's definition is less operational than de Finetti: depending on the understanding of "dispositions", dispositions to act do not correspond to any actually measured values, but only say that an agent would act in such and such a way, given certain things applied. Of course, actions are determined by many other factors then just beliefs. Ramsey states his extra and vital assumption needed for measurement quite clearly: "I propose to take as a general psychological theory [...] that we act in the way we think most likely to realize the objects of our desires, so that a person's actions are completely determined by his desires and opinions." (Ramsey 1926, 35) Given this, rational actions are what an agent *believes* to most satisfy his *values*. Mathematically, rational choices to act are functions from utility and probability. Given that we try to obtain what we want most—that we can be represented as maximizing utility—in a measurement scenario we can derive (if coherent) a probability function. Thus, actions are taken as fully determined by degrees of belief and values of the agent; probability is defined as rational *acts*, be that acts of betting or forecasting (or some other measurement scenario). Whether this definition is plausible we will consider below, after briefly discussing whether the assumptions made about preferences in defining subjective probability are convincing.

3.1.2 Interval-valued degrees of belief

As Walley (1991, 243) observes, it is not clear that everybody has *fair* prices, meaning that everyone considers in every situation of uncertainty what certain gain is equivalent to the random gain. Thus, it is not clear that everyone has one precise price for a random gain or gamble, rather than an interval of prices; an agent might be indifferent between several exchanges for gambles. If prices are supposed to be an agent's prices for bets, where the prices are defined as preferences over gambles (or the desirability of gambles) as Walley does (Walley 1991, 60–66), forced two-sided elicitation does not necessarily measure an agent's prices accurately. An elicitation of fair prices assumes a too strong theory of decision making (Walley 1991, 242–243).

For this reason, the theory of imprecise probability as Walley constructs it does not define probabilities by means of fair prices, but allows buying and selling prices for gambles to differ, yielding a theory of imprecise probabilities and previsions. The mathematical details of this approach are much more complicated and will not be discussed here; the important thing here is what the theory does and what is does not solve. It allows for difference of buying and selling prices, such that the agent need not have precise preferences over exchanges between random and certain gains: she can be unsure, and noncommittal to some degree (interval). This leads to a "one-sided" betting scenario, as opposed to the two-sided betting scenario of de Finetti. The one-sided betting scenario is taken by Walley as a better definition of subjective probability, because it requires a less strict decision-theoretic background. The prices—now not fair—are still defined by means of a preference relation over random gains, with the difference that the preference is not defined by exchanges with any certain gain. Subjective probabilities are still rational dispositions to act.

De Finetti realized that his definition of subjective probability by means of fair prices was forced and an idealization, but said that his theory of precise probabilities is an idealization no different from idealizations in many other empirical sciences and branches of applied mathematics (Vicig and Seidenfeld 2012). *Ideally* anyone would (if rational) have coherent fair prices, given enough contemplation of specific problems. The theory of subjective probability is a mathematical idealization of agents' dispositions to act, and the problem of unclear preferences and thus imprecise probabilities was taken by de Finetti as a problem of incomplete elicitation, foremost.

The question of whether prices are said to be interval-valued or precise does not affect our discussion of Epistemic Puritanism: on both construals, probabilities by definition involve values and preference. Some Epistemic Puritanists have said that their assumption that beliefs come in degrees is an idealization, and that degrees of belief should be taken to be interval-valued as in the theory of Walley. For example, Joyce writes: "Most probabilists recognize that opinions are often too vague to be pinned down in numerical terms, and it has therefore become standard to represent a person's partial beliefs not by some single credence function but by the class of all credence functions consistent with her opinions" (Joyce 1998, 600).[2] However, if there is no clear way to assign precise numbers to beliefs because of the lack of a measurement scale, there is also no clear way to assign interval-valued degrees to beliefs. Joyce also neglects explaining what the "consistency" of a group of credence functions—credence functions are a superset of the set of price functions, where credence functions can also be incoherent—with beliefs, is. If the notion of numerical pure degrees of belief lacks meaning, my considering rain in April 2015 in Munich unlikely is perfectly well consistent with a number 0.9, or an interval of 0.85 – 0.95. A measurement scale is not made redundant by resorting to interval-valued degrees of belief; a measurement scale is just as important here, to specify the meaning of the numbers. Compare this with saying that the temperature is 50; never mind the scale.

This is not to say that the theory of interval-valued degrees of belief is of no importance to Bayesian Epistemology. If values *are* allowed in the setup, agents with specific opinions

2. Christensen, whom we will discuss later, is of a similar view (see footnote 1 of chapter 4). See also the collection of essays in (Huber and Schmidt-Petri 2009), where most philosophers assume that beliefs do come in degrees, perhaps imprecise degrees.

can be said to be willing to bet on a range of prices; this is the theory of Walley. But our concern in this thesis is Epistemic Puritanism, where values are *not* allowed in the setup; in the case of Epistemic Puritanism, a measurement scale must be provided because a range of betting prices still involves values and acts. The question of the meaning of degrees of belief is therefore the same for the theory of precise and imprecise probabilities in the case of Epistemic Puritanism; for simplicity we will not discuss imprecise probabilities further here.

3.2 Measurement and dispositions to act

On the interpretation that we constructed in the previous chapter, probability is a representation of rational dispositions to act in situations of uncertainty. This is a more general interpretation of de Finetti's setup and less strictly operational than de Finetti's own interpretation, but it is still operational or behavioural in that it defines degrees of belief as possible acts. They are dispositions to accept betting contracts (buy or sell) or dispositions to forecast, given that penalties are attached to the values that are elicited. Before being able to see if this can be interpreted or adapted according to the demands of Epistemic Puritanism and whether this behavioural or operational definition is plausible, we must first get clear on what is meant here by a disposition to act. Consider for this a basic everyday example: a road cyclist who has just climbed the *Passo dello Stelvio*, is taking a short break and will have to take the descent in a few minutes. It has just started to drizzle, but stronger rain was forecasted the day before. We can see the cyclist pondering on how many clothes to put on for the descent; suppose we know the route he is going to take and can see dark clouds in that direction. If he considers rain and a drop of temperature more likely, he will be disposed to put more clothes on; if he thinks he will be fast enough to be down before the rain and believes it will be warm enough, less. Such considerations, of course, we could ask him directly, but these are not clearly and non-ambiguously numerical, and therefore cannot be represented by a probability function. That is, without further testing so as to see whether the axioms of chapter 1 are satisfied. But if we were to subject him to a measurement scenario, this could be obtained.

When we say that an agent *would* elicit such and such values in a specific measurement scenario, this is not to say that we consider the agent's opinions and beliefs to be already numerical and that these degrees of belief are measured by a scenario—numbers already in the head and only extracted, so to speak. That is not how the numbers of the setup of the previous chapter are assigned; the numbers are assigned by measurement. In a measurement scenario, an agent is forced to post a single value (be that a price or a forecast) and it is in the agent's interest that this single value is also the agent's best estimation: there are penalties attached to this value. To post such a value is an *act*; probabilities are in effect defined as possible acts, either acts to accept betting contracts or acts to forecast in a scoring rule scenario. These acts in the measurement scenarios are numerical, but that is *not* to say that the beliefs are also already numerical.

The betting scenario builds on the decision-theoretic foundation of prices, and assigns a number to an opinion of an agent through the willingness to exchange money for gambles. Because of the problems with the betting scenario as an elicitation scenario, this can be taken as a willingness to act in ideal conditions; conditions where an opponent does not distort the opinions, where there is no risk-aversion or risk-seeking, etc. The scoring rule scenario ascribes numbers to opinions through the usage of a strictly proper scoring rule. If another scoring rule were used that is not *proper*, other values could be elicited rationally and a definition of probability as the elicited value fails: opinions are not accurately elicited. Similarly, is it is not *strictly proper*, there is no single value. The definition of probability in the scoring rule scenario relies on the penalty for the agent; the definition is in terms of possibly measured values. So, both scenarios can be interpreted as being about dispositions to act. However, an identification of beliefs with dispositions to act is neither plausible nor necessary, as I shall argue.

3.2.1 Dispositions to act and beliefs

A full treatment of the concept of disposition extends far beyond what we can discuss here and is also not necessary for the purposes of this thesis. What is important is that a disposition to act cannot generally be identified with (folk-psychological) beliefs. Many beliefs do not lead to action and although all beliefs can perhaps be understood as *possibly* leading to some action—that every belief can be understood as a disposition to act in appropriate circumstances—beliefs cannot be strictly connected with any actions. I might love coffee in the morning very much and will therefore generally be disposed to make coffee in the morning, but there are always conditions imaginable under which I will not make coffee. This we have also already discussed in chapter 1. This argument also extends to the betting scenario or forecasting scenario: specific opinions of likelihood do not always lead to specific betting contracts nor to forecasts. If all problems of measurement discussed in the end of the previous chapter are dealt with, we can infer the agent's considerations of likelihood, but this is not to say that these considerations of likelihood *are* the elicited values nor that these are numerical. They are numerical only in combination with values: it is the measured value that is numerical.

Dispositions to act are possible acts, and therefore also include values. Even if problems of measurement make it so that any measurement is slightly distorted, *ideal* measurement is used for the definition of probability. An example of such a general problem of measurement is that probabilities of multiple decimals are certainly possible (a probability of 0.35783) such probabilities would probably never be elicited. But also in such an ideal interpretation of dispositions to act, precise numerical degrees are ascribed by this ideal measurement scenario: if measured correctly, the agent would post such and such values. That is what a degree of belief as a disposition to act *means*. This ideal measurement can be taken to mean that an agent would post such values given enough contemplation; the problems are then problems of incomplete elicitation, as was also de Finetti's remark with respect to imprecise probabilities, as seen above.

This resorting to ideal situations is not possible for the epistemic puritanist. It is not clear how beliefs are to be measured in a purely epistemic way, not even ideally. Beliefs only lead to actions if the action is also deemed desirable for the agent; if the agent has no values, rational actions cannot be said to exist. Epistemic Puritanists try to justify subjective probability theory through one of the measurement scenario's as we shall see in the next chapter, but cannot use the definition of degree of belief—betting prices or forecasts—that is used in the setup of de Finetti.

Even if a purely theoretical definition of probability as prices is adopted as in the setup of **P** of the previous chapter, with no empirical import and way of measuring, the numerically precise value of this function and what it represents is still defined through an *act*: an exchange of a certain gain for a random gain or gamble.

As seen, Ramsey explicitly defined degrees of belief as dispositions to act. This is not to say that Ramsey thought beliefs alone cause actions, or that beliefs themselves are measurable. What Ramsey does is redefine the concept of belief in behaviouristic (operational) terms: this can be considered *methodological* behaviourism (Sellars 1956).[3]

3. I am not here explaining Sellars' complete theory nor using it for the same goals as he: he uses his arguments to the effect that our folk-psychological notions are based on behaviour. I am adapting part of the argument to apply to our setup of the previous chapter. Folk psychology remains in Sellars' theory; his view was along the lines of Davidson (Davidson 1980b, 1980c) that the intentional and referential is indispensible in accounting for human behaviour and beliefs (deVries 2015). I am sympathetic to this view, though it is not of importance to the arguments made in this thesis. I am not claiming here that folk

This sort of behaviourism is not an attempt at defining folk-psychological concepts in terms of only observable behaviour.

Choosing to define degrees of belief as dispositions to act is redefining the notion of belief and in doing so, disregarding some aspects of the old concept. Jeffrey has noted the same thing: "[I am not] disturbed by the fact that our ordinary notion of *belief* is only vestigially present in the notion of degree of belief. I am inclined to think Ramsey sucked the marrow out of the ordinary notion, and used it to nourish a more adequate view."(Jeffrey 1970, 172) Ramsey's notion of degree of belief should not be understood as an understanding of psychological beliefs; rather, they are dispositions to act. Dispositions to act are taken as primitive and whatever connection is there with traditional (folk-psychological) beliefs, is only of secondary interest.

The conclusion of the argument is that the setup does not work with pure beliefs: we do not have a way of assigning numbers to pure beliefs, because we have no way of numerical measurement. Epistemic Puritanism relies on pure beliefs, and wants a definition of subjective probability as being about only beliefs; because beliefs are not strictly connected to values and acts, the definition cannot use values and preferences. But because of the operational definition of Ramsey and de Finetti, we have no definition of degree of belief that the Epistemic Puritanist can use. The measurement scenarios used for the meaning of the numbers—the scenarios provide a measurement scale—cannot be adapted for the purely epistemic case. It is therefore unclear what epistemically pure degrees of belief are; let alone what coherent epistemically pure degrees of belief are, which are to be represented by a probability function.

3.2.2 Dispositions to act and (anti)realism

How does this relate to the question of realism, as we have discussed it in section 1.3? Zynda's view was a weak-realist view with regard to probability: a theory in which only considerations of likelihood are said to exist, and where probability is defined by logical construction on these considerations of likelihood. Applied to the theory just discussed, I consider such a realist view of dispositions to act undesirable and unnecessary: dispositions to act are idealizations about agents. It is not assumed that agents have beliefs that actually lead to betting prices or forecasts nor that agents are actually measured at some point and time; betting prices or forecasts are taken as primitive, but these are not psychological

psychology can be replaced by any other framework or theory, only that subjective probability theory is not part of folk psychology.

things as beliefs are. The justification for talking about dispositions to act does not come from their existence or their truth of the matter, but from the pragmatic virtues of the theory of which they are part: decision theory. This theory is a strong theoretical and psychological idealization.

The purely epistemic approach can be viewed as a strong realist interpretation of probability and degrees of belief. Degrees of belief are taken to be real psychological things, and if these numerical degrees of belief are coherent, they *are* probabilities. This is not to say that the epistemic puritanist need be a strong realist, but realism does seem to be a natural position for the epistemic puritanist given that probability functions are defined (as in the setup of the previous chapter) according to the two coherence conditions of boundedness and additivity. Such a strong realist position is problematic, as we have seen: we do not have a clear notion of what it would mean to have beliefs that come in numerical degrees. We do not have a clear way of making sense of the numbers. The weak realist position is more plausible, where we would have to find a way to construct a general, not strictly operational, theory of subjective probability either from psychological entities or from measured values. Here probability can be said to represent these measured values or opinions. Of course, defining subjective probability by measured values is certainly possible—this is the strictly operational interpretation de Finetti seems to have held—but such an interpretation is not very general. A weak realist account based on dispositions is broader, but more difficult: dispositions to act are idealizations about agents, which can be read to say that an agent would act in a specific way if certain conditions held. Dispositions to act involve many ceteris paribus clauses and are idealizations about agents.[4] A weak-realist view would have to construct dispositions to act from logical construction from something else; I am not sure if this can be done. If degrees of belief are taken to precise, a weak realist view of probability is possible by saying that probability is a representation of these degrees of belief; but as we have argued, such a notion of degrees of belief is ambiguous.

An antirealist position, where dispositions to act are used in virtue of their pragmatic virtues, seems the best interpretation with regard to a general theory of subjective probability. Here probabilities are defined as representations of dispositions to act, but these dispositions to act are not said to be "real" psychological attitudes. We use these in explaining behaviour and as idealizations about agents' beliefs, and justify the usage of the

4. See (Levi and Morgenbesser 1964) for an investigation into the notion of dispositions to act, which identifies them as theoretical idealizations about agents and their beliefs, idealizations which can be seen to be about beliefs together with many ceteris paribus clauses.

notion on pragmatic grounds. Although such antirealism is possible for the epistemic puritanist, it does not seem the natural position: it would be to say that degrees of belief are not real attitudes. The normal way for an antirealist to justify talk of degrees of belief is on pragmatic grounds, but this is not plausible for the epistemic puritanist: this would require, in the end, justifying epistemic puritanist Bayesian Epistemology on the basis of pragmatic virtues (explanatory strength, simplicity, etc.). These pragmatic virtues do not seem to be purely epistemic, and therefore antirealism does not seem tenable for the epistemic puritanist.

3.3 Criticisms from epistemic puritanists

I have explained how I think the theory should be (and how it should not be) read. However, many critics have not found the setup and construction of de Finetti and Ramsey convincing or plausible.[5] Many of these criticisms come from the standpoint of Epistemic Puritanism: the construction of subjective probability is not a theory of (degrees of) belief, even though it should be. They desire a purely epistemic construction of subjective probability, and want to get rid of the values inherent in dispositions to act. Let us now consider these criticisms.

One philosopher whose criticism is clear and provides a good summary is David Christensen. Christensen finds the definition of degrees of belief by means of preferences and acts unsatisfying. Relating to remarks about the definition of degrees of beliefs as acts, he writes: "the move of defining degrees of belief in terms of an agent's preferences (as revealed in her choice-behaviour) is reminiscent of the standard operationalist strategy in philosophy of science: taking one way of measuring a theoretical quantity and treating it as a definition."(Christensen 2004, 108) Christensen focusses on the betting scenario and the Dutch Book argument, and objects to the betting interpretation that it requires a too strong connection between beliefs and betting prices. He says that "[an] acceptable interpretation of the [Dutch Book arguments] must acknowledge that partial beliefs may, and undoubtedly do, sometimes fail to give rise to the preferences with which they are ideally associated"(Christensen 2004, 113–114). The Dutch Book argument is seem by him to be

5. Some of these critics who have tried giving constructive purely epistemic adaptations to the setup we will discuss individually in the next chapter. Some notable critics are Kyburg (Kyburg 1978), Christensen (Christensen 2004), Joyce and other accuracy-first epistemologists (Joyce 1998; Leitgeb and Pettigrew 2010a; Pettigrew 2013), Hájek (Hájek 2008; Eriksson and Hájek 2007).

intended to indicate epistemic irrationality, whereas—because of the strong link between beliefs and betting behaviour implicit in the betting scenario—it actually only establishes pragmatic irrationality.

We will come back to Christensen's remarks on epistemic and pragmatic rationality in the next chapter. For now, let us focus on the connection between beliefs and behaviour; he finds this operational definition implausible and the source of the troubles with the Dutch Book argument. It is however not clear that this definition is problematic in this case. Beliefs are only redefined for the purposes of subjective probability theory, where it is required for the theory to go through. This is not to say that we have to assume an eliminativist position to folk-psychological beliefs: subjective probability theory is perfectly compatible with a theory that says that folk-psychology is indispensable. The concept of belief is just not used in the construction of subjective probability. To be more precise, what is (re)defined is the notion "degree of belief"; but this was not a notion that already existed—not in a clear case, anyway—in folk-psychology. The old concept of belief need not be disposed with; it is just not useful for the purposes of decision theory and subjective probability theory. Furthermore, de Finetti's lack of usage of the notion of degree of belief also indicates that subjective probability theory is not a purely epistemic theory at all, nor intended in that way.

Related criticism have been made by other philosophers, also usually focussed on the betting interpretation. Here are some of the criticisms summed up: Eriksson and Hájek[6] argue against the setup with the example of a Buddhist monk with no desires and values—no substantial ones anyhow, he presumably does not really care about money. Surely this monk has beliefs about some event being the case and surely (they say) he has degrees of belief, but he has no (betting) price (Eriksson and Hájek 2007, 194). Similar remarks have been made by Christensen, stating that the relation between beliefs and preferences in the setup and argument is undesirable (Christensen 2004, 106–115). As regards especially the penalties in the measurement scenarios which are also given in some measure of linear utility, James Joyce has made similar points, focussing on the central role in the setup given above on maximizing expected utility and wanting to minimize expected loss, rather than on purely *epistemic* merits; a penalty in something purely epistemic. Kyburg has argued that an agent would have to be *forced* to bet, but that avoiding sure loss in forced situations

6. It is interesting to note that Eriksson and Hájek at the end of their paper propose taking degrees of belief as uninterpreted primitives. This is a position somewhat like the one I have explained in this chapter, with the difference that a completely *uninterpreted* primitive is certainly undesirable.

has nothing to do with epistemic rationality, only pragmatic rationality (Kyburg 1978). Lastly, it has been said that the irrationality of incoherence is not established properly. There might not be any clever bookies in the world who can make me elicit (firstly) my degrees of belief accurately, but more importantly can then stipulate the stakes such that I would lose no matter what happens. In fact, if there are no such smart bookies but only rather foolish ones, it might even be *rational* to be incoherent, because there exists a set of bets for any incoherent set of values for which the gain is uniformly *positive* (Hájek 2008).

It is first important to note the emphasis placed on the betting scenario in these criticisms, where the scoring rule argument avoids some of the problems. The requirement of a smart opponent is not required in the scoring rule setup, and that part of the criticism is therefore only a particular problem of one way of measurement. Aside from that, the idea that there have to be bookies running around is simply a misunderstanding of the theory: it is about dispositions to act, where the betting scenario is an ideal measurement scenario. But more important in most criticisms is the emphasis on *degrees of belief* as not being identical with betting prices. If there *were* degrees of belief, this would be true; the problem is that it is not clear that there are degrees of belief in the purely epistemic sense. And if the theory is understood as being about dispositions to act, especially in the scoring rule scenario, the problems raised about the existence of the appropriate smart bookies disappear.

Before closing off this chapter, it is interesting to note another interpretation which does not take De Finetti's setup as being about dispositions to act, and which does away with the emphasis on preferences and values. This is the interpretation advanced by Howson, who also takes beliefs out of the picture: he proposes to understand the coherence conditions—more precisely, the properties of **P**—as similar to logical consistency conditions and thus not in need of further explication. Coherence can already be defined in the setup, as properties of **P**—boundedness and additivity—without bringing in any kind of (practical) irrationality or avoidance of sure loss. According to Howson, the coherence conditions can be understood as conditions for *consistency*; the theory of subjective probability is very much like a many-valued logic. Howson points to some remarks of De Finetti, especially his earlier work, that hint to such a conception.[7] The problem with this approach here is that

7. For example, "it is better to speak of coherence (consistency) of probability evaluations rather than of individuals, not only to avoid this charge, but because the notion belongs strictly to the evaluations and only indirectly to the individuals"(de Finetti 1980, 63), and "[it] is beyond doubt that probability theory can be considered as a multi-valued logic (precisely: with a continuous range of values), and that this point

the theory is no longer a normative theory about *beliefs*, either.[8] It therefore does not satisfy the conditions of a purely epistemic interpretation of subjective probability: we wanted a theory of rational beliefs, and Howson's interpretation does not treat subjective probability theory as about beliefs. Also, De Finetti's usage of the Dutch-book argument and scoring rules do point to such a normative interpretation of beliefs or decisions, as Howson also recognizes (Howson 2009, 114). So we do want an argument for these coherence conditions, relating them to beliefs (or decisions) in a normative way.

Kyburg, an influential critic of subjective probability theory and Bayesian Epistemology, draws the conclusion that there are no psychological entities which correspond to degrees of belief (Kyburg 1978, 175). I have here argued for the same position.[9] We can here draw a similar conclusion: it is not clear what purely epistemic degrees of belief, that can function in the same way as the prices or forecasts from De Finetti's setup, are or could be. There is no purely epistemic measurement scale by which we can meaningfully and nonambiguously assign numbers to beliefs. Degrees of belief in the practical sense of Ramsey and De Finetti are defined in terms of personal values: in the betting scenario most clearly by being defined by preferences over acts (exchanges), in the scoring rule scenario because forecasts are subject to a penalty. Take away the penalty or the values and preferences, and you take away the definition of a degree of belief. Subjective probability theory takes *acts* as primitive, not purely epistemic beliefs. Furthermore, trying to understand purely epistemic degrees of belief as dispositions to act in the way required is precluded by the very notion of belief; only through rigorous redefinition can the setup of subjective probability go through.

of view is the most suitably to clarify the foundational aspects of the notion and the logic of probability." (De Finetti in Howson 2008, 4). Another important argument for Howson's construal of the theory comes from De Finetti's postulation of finite additivity and rigidity, which has the form of an axiom in De Finetti's theory; we shall not go into this topic here. See (Howson 2008, 2009).

8. Cf. footnote 17 of chapter 1.

9. However, Kyburg also thinks subjective probability theory construed as about decisions and dispositions to act is "vacuous" (aside from the general theory being "philosophically bankrupt")(Kyburg 1978, 179), because it does not give any *guidance* as to how to act—something he considers a requirement of any normative theory of decision making. I do not agree with this criticism, although this extends beyond the purposes of this thesis. I believe Bayesian theories constitute, in Blackburn's terms, a "grid for imposing interpretation: a mathematical structure, designed to render processes of deliberation mathematically tractable, whatever those processes are" (Blackburn in Peterson 2009, 454).

Chapter 4

A Purely Epistemic Justification?

Let us suppose that through some scientific breakthrough (and psychological miracle), it is discovered that there are, in fact, purely epistemic degrees of belief, and that everything said in the previous chapter on the topic is false. Starting with degrees of belief, we are still left the task of justifying the coherence conditions to obtain a probability representation; as we have seen, these conditions too required, in de Finetti's construal, preferences. It is in an agent's interest to minimize expected loss. Can an epistemological justification be given that fits better with the epistemological approach desired of Bayesian Epistemology?

4.1 Adapting the Dutch Book argument

Some have attempted interpreted the Dutch book argument as showing an epistemic defect. An elaborate and recent account of this can be found in the work of Christensen, whom we have already seen in the previous chapter. Christensen aims to get rid of the connection between beliefs and values in the argument and replace the *pragmatic* irrationality in the Dutch book argument—being incoherent leaves you victim to the possibility of a Dutch book being made against you—by *epistemic* irrationality. Christensen sees the Dutch book argument as showing an inconsistency in an agent's values and preference, if an agent is incoherent: he views the coherence conditions on prices $\mathbf{P}$ as consistency conditions on preferences, and if an agent's betting prices do not follow these coherence conditions the agent is inconsistent (Christensen 2004, 111–113). But this is only pragmatic inconsistency, whereas Christensen desires the Dutch Book argument to show epistemic inconsistency: there is something epistemically wrong with being inconsistent in your degrees of belief.

In rejecting the interpretation of Dutch Book arguments and in taking beliefs as primitive, he rejects defining degrees of belief as betting prices or as dispositions to act; his emphasis on being interested in inconsistency of purely and only beliefs justifies considering his attempt an epistemic puritanist attempt at justifying the coherence conditions.[1]

His account is based on the "intuitive" idea that a specific degree of belief *sanctions as fair* a specific betting price. A particular degree of belief provides a "justification for the agent's bet evaluation—it is part of what makes the bet evaluation a reasonable one"(Christensen 2004, 116). This justification means that a bet placed on an event A at, say, 2:1 odds, is justified by an agent considering the event A to be twice as likely to happen as not to happen.

The primary problem that Christensen then turns to is that of "value interferences" (Christensen 2004, 117). He explains value interferences as that an agent's preferences or values at particular situations might interfere with the betting prices and therefore not accord to the agent's degrees of belief. Thus, having other outstanding bets at a time[2] might interfere with some other bet. Plus, problems of the non-linearity of monetary utility also interfere with betting prices. Christensen wants the betting setup to be so formulated as to remove these possible value interferences. For this purpose, he formulates his "depragmatized" Dutch Book argument in terms of "simple agents": for simple agents, there are no value interferences. And because there are no value interferences for simple agents, simple agents *should* evaluate bets as being linked directly to his degrees of belief; a simple agent's degrees of belief *sanction as fair* bets placed at odds matching his degrees of belief (Christensen 2004, 117).

The irrationality is then explained also in terms of the betting prices of simple agents: if a simple agent has degrees of belief that sanction certain betting prices as fair that are incoherent, that simple agent sanctions as fair betting contracts that would leave him

1. Interestingly enough, Christensen acknowledges the importance of measurement in defining the notion of degree of belief (Christensen 2004, 114) and acknowledges that degrees of belief cannot be assumed to be precise, but does ascribe to a view of degrees of belief as being imprecise, in the sense of interval-valued numerical degrees of belief (Christensen 2004, 143–149). In the light of the argument of the previous chapter, I consider his conception of degrees of belief unclear. It is also unclear in how far Christensen adheres to Epistemic Puritanism: many of his comments hint to Epistemic Puritanism, but his theory (as I shall explain) does not seem to be purely epistemic.

2. Although this might sounds like a diachronic problem, it should be read as a synchronic one. Problems of bets placed at different times are not meant, *as* placed previously, to interfere with an agent's current betting prices. The argument is meant only to indicate that other values can interfere with an agent's betting prices.

worse off no matter what happened. This is irrational. To then return to the case of normal agents, Christensen contends that the only difference between normal agents and simple agents are the value interferences. Thus, incoherent degrees of belief are degrees of belief that sanction as fair betting prices in betting scenarios free from value interferences, where these betting prices are such that a Dutch Book could be made against the agent on the basis of these betting prices. By contrast, coherent degrees of belief sanction as fair betting prices on the basis of which no Dutch Book could be made against the agent, in scenarios free from value interferences.

This explanation seems like a plausible explanation, but it is too much like the idealized dispositional explanation of the betting scenario we have given in the previous two chapters to be a purely epistemic explanation. The problem of value interferences is one of the main reasons why de Finetti went on to prefer the scoring rule scenario as opposed to the betting scenario; value interferences are, if at all, much less present in the scoring rule elicitation scenario. But this scenario is still not purely epistemic, because values and preferences are still essential part of the measurement: both in the measurement and in the justification of the coherence conditions. And Christensen's explanation of the Dutch Book argument also still involves values and preferences in an essential way: the irrationality is still that a series of bets can be made against an agent that would leave him *monetarily* worse off—this hardly counts as purely epistemic irrationality. And getting rid of pragmatic irrationality is exactly what Christensen desires and intends his account to accomplish: he wants the Dutch book argument to show epistemic irrationality.

Christensen, though, thinks this is accomplished in his account because the fact that degrees of belief *sanction as fair* specific betting odds makes the property of being open to a Dutch Book a property of the degrees of belief, not of the betting prices. I see two problems with this: firstly, why an agent should not be incoherent in her degrees of belief, regardless of what they sanctions as fair; secondly, what this "sanctions as fair" relation is that turns pragmatic (ir)rationality into epistemic (ir)rationality.

The first point is that the irrationality of the setup and construction of de Finetti makes it undesirable to be incoherent. On Christensen's account, irrationality of degrees of belief is that betting prices sanctioned as fair by degrees of belief lead to sure loss; but why any agent should care about what degrees of belief sanction as fair is unclear. There is no need to worry at any point—not even ideally or hypothetically—that whatever my degrees of belief sanction as fair could lead to financial disaster; my degrees of belief themselves do not lead to financial disaster, and if I ever have to bet I can just post coherent bets.

However, this criticism can be said from the purely epistemic standpoint to be question begging, insofar as it demands rationality that is non-epistemic and undesirable from the agent's perspective.

But the criticism is linked to the second, more fundamental, problem: the sanctioning as fair relation is very ambiguous. Irrationality is cashed out in terms of possible *acts*—buying and selling bets—and the sanctioning as fair relation is supposed to somehow transfer this to purely epistemic irrationality. If values are taken out of the picture, there is no irrationality, practically or epistemically, because irrationality is still cashed out in terms of a Dutch Book; because Epistemic Puritanism takes values out of the picture, irrationality is taken out of the picture. Christensen's "sanctions as fair" relation is all the more surprising in light of his rejection of what he calls the "metaphysical" relation between degrees and belief and preferences (Christensen 2004, 129). If anything, the relation he offers is much more metaphysical than the definition of degrees of belief as dispositions to act we have seen in the previous chapters. Degrees of belief, on Christensen's account, are strongly connected with betting prices (but not defined as betting prices), and incoherence of degrees of belief is irrationality of possible betting prices and thus practical irrationality.

I conclude that Christensen's account does not succeed in defining purely epistemic irrationality: values are still essentially required for the definition of irrationality. His attempt at understanding incoherence (or inconsistency) of preferences and values as epistemic defects fails: subjective values are not purely epistemic, yet essentially required for the Dutch Book argument. An account that might seem hopeful at this point is to replace the pragmatic subjective values by purely epistemic values: this has been attempted with regard to the forecasting scenario.

4.2 Adapting the Scoring Rule argument

Where critics of subjective probability theory have focussed mainly on the Dutch book argument, proponents of an epistemological account have lately celebrated mainly the scoring rule setup. Adapting scoring rules as highlighting epistemic flaws was first proposed in a detailed manner by James Joyce (1998), and has been picked up and further constructed by many formal epistemologists (Pettigrew 2013; Leitgeb and Pettigrew 2010a, 2010b; Greaves and Wallace 2006).[3] This version of Epistemic Puritanism uses scoring rules as

3. In our discussion we will focus mainly on Joyce's work, although all of the arguments given in this chapter are general and do not depend on any particular aspects of his approach. Constructions of accuracy-first

measures of accuracy of degrees of belief, and can therefore be called *accuracy-first epistemology*. The idea is surprisingly simple, given the mathematical and technical work behind it: accuracy-first epistemology takes scoring rules to measure inaccuracy, which they conceive as an epistemic defect. The coherence conditions of degrees of belief are then justified by the argument that incoherent degrees of belief are strictly less accurate then some other set of degrees of belief.

Scoring rules on this approach are taken as measures of accuracy of degrees of belief, where what was the penalty in the approach of chapter 2 is now a measure of *in*accuracy: the higher the score, the less accurate the degree of belief. Accuracy is understood to be an epistemological norm, similar to the truth norm of full belief; degrees of belief *ought* to be as accurate as possible. More precisely, Joyce formulates the norm as a norm of *gradational* accuracy (Joyce 1998, 579): agents must maximize their degree of confidence in true propositions while minimizing their degree of confidence in false propositions (Joyce 1998, 578). So although maximal accuracy is obtained by always being spot on about all propositions—believing all truths to degree 1 and all falsehoods to degree 0—whenever an agent is uncertain she should be encouraged, by the scoring rule, to be as accurate as she can be; avoid being confident in falsehoods but be as confident as possible in truths. The analogue in accuracy-first epistemology of de Finetti's reliance of minimizing expected loss is to minimize inaccuracy.

Although the formulation is of scoring rules as "measures" of inaccuracy, a fundamental difference with de Finetti's usage is that they cannot be taken as measurement or elicitation scenarios of degrees of belief. Let us briefly for this paragraph put our assumption that there are degrees of belief between brackets, and consider the relation between the accuracy-first epistemologists usage of scoring rules and de Finetti's usage. The idea in the measurement scenarios of de Finetti that we saw in chapter 2 is that an agent is *penalized* in some measure of value and that it is, for that reason, in the agent's best interest to elicit her best estimate. To say that the penalty is epistemic and that it is epistemically wrong to be further away from the truth than you have to, does not make it in an agent's best interest to post her best estimate. Whereas in de Finetti's theory forecasts are made precise in virtue of being subject to the penalty—strictly proper scoring rules force a *single-valued* forecast, that is also the agent's *best estimate*, through penalties that are of value to the agent—this cannot be done in the purely epistemic case. A measure of inaccuracy does not make it in

epistemology differ in the exact characterization and definition of dominance and in their exact understanding of epistemic utility, but all use the notion of accuracy and use this in their concept of epistemic utility.

the agent's interest and therefore does not work as an elicitation scenario in the same way. Supposing that the right scoring rule for the purposes of measuring epistemic accuracy is the Brier score, and that the agent is forced to elicit one value, there is still no incentive for the agent to elicit his best estimate. The agent need not *value* epistemic values. The setup of accuracy-measures is therefore to assume that there are degrees of belief, which are held to some measure of gradational accuracy and are penalized in the pure epistemic sense of accuracy. This is also noted and highlighted by Joyce, who says that degrees of beliefs are *not* actions and are not *chosen* and scoring rules attach to degrees of belief ("credences" in his terminology) directly, not to measured values (Joyce 2009, 266, footnote 5).

Scoring rules, on this purely epistemic approach, are then used as measures of epistemic utility, where gradational accuracy is taken as the ultimate measure of epistemic utility. We will return to this notion of epistemic utility below. Note, though, that we cannot here simply adopt the Brier score and say this measures inaccuracy or epistemic (dis)utility: we have to argue that this is a good measure of inaccuracy. This is also unlike de Finetti's measurement scenario and setup, where values are measured such that it is in the agent's interest to *minimize expected loss*, and the Brier score then makes it in the agent's interest to elicit her precise best estimate of some random variable. The Brier score allows us to get what we want from the agent. But because scoring rules measure accuracy, we have to be sure that a specific scoring rule is an accurate gauge of (in)accuracy. In (Joyce 2009; Leitgeb and Pettigrew 2010a) it has been argued from other epistemic principles that the Brier score is this best measure based on other principles about epistemic values, which might be contested but will not be discussed further here. But whatever scoring rule will turn out to be the one that best measures accuracy, the scoring rule can be represented by a function $\mathbf{I}(\mathbf{b}, w)$, where $\mathbf{b}$ is a degree of belief (not necessarily coherent), and w is a set of truth value assignments (a "possible world"). The argument for probabilism—that an agent's degrees of belief should satisfy the properties of $\mathbf{P}$ listed in chapter 2—is then made by saying that any coherent set of degrees of belief is *undominated* by any other set of degrees of belief, whereas any incoherent set *is* dominated by some set of degrees of belief. A set of degrees of belief $\mathbf{b}$ is said to be dominated by another set of degrees of belief $\mathbf{b}'$ if and only if $\mathbf{I}(\mathbf{b}, w) > \mathbf{I}(\mathbf{b}', w)$ for every truth-value assignment w (Joyce 2009, 267). This says that an agent would be more accurate *whatever state of the the world* obtained, or would be more accurate in every possible world.

Criticisms of this account have mostly focussed on whether accuracy can indeed be taken as a cardinal epistemic virtue and norm (Easwaran and Fitelson 2013), whether the epistemic principles used as a background for scoring rules measuring epistemic utility are plausible (Maher 2002), and whether it is really established that coherent degrees of belief are never dominated, whereas incoherent degrees of belief always are (Hájek 2008). What is more interesting and problematic is what is actually measured; what is this distance from the truth which the scores are supposed to measure?

4.2.1 Epistemic Utility

Howson has, in this light, questioned whether it is actually accuracy with respect to the *truth* that is measured: the number 1 as proxy for truth is only convention and could be switched around, with the result that every coherent set of degrees of belief is dominated by an incoherent set (Howson 2008, 20). I find Howson's short remark on the matter ambiguous, though what I take this worry to mean is that distance with respect to *truth* cannot be assured by measuring the distance of a credence to 1. However, as Joyce argues (Joyce 2009, 264), if truth values are switched around the accuracy toward the truth would simply be the distance to 0, rather than 1, and degrees of belief also have to adapted accordingly (by substracting it from one). Alternatively, it seems to me that a viable option here is to define an *ideal* set of degrees of belief that is right about every proposition at a world, as is Pettigrew's approach, and avoid the problem by stipulating that truth is 1 in this way (Pettigrew 2013, 899). But this criticism of Howson leads us to the question of what accuracy measures; that is, whether accuracy is actually a measure of something called epistemic utility. For it is unclear that there *is* a numerical measure of epistemic (dis)utility that can be based upon (in)accuracy.

Such is the criticism of Mayo-Wilson and Wheeler, whose argumentation I will here summarize (Mayo-Wilson and Wheeler 2015). The fundamental problem is that scoring rules yield numbers quantifying loss in some measure of linear utility, but construing purely epistemic accuracy as a measure of utility seems impossible. To start off, it is not clear from the get go that all beliefs can indeed be strictly compared in accuracy: it is not clear that my degree of belief of 0.7 that strawberries in the supermarket in June are good is more accurate than my degree of belief of 0.3 that there would be snow to come in April 2015 in Munich. If it is, this is not to say that accuracy is also numerical and, moreover, that all beliefs can be ascribed the same measure of accuracy by which they can be compared. Joyce here recognizes the assumption that epistemic utility can be quantified as a fiction—albeit a "useful fiction" (Joyce 2009, 266). He defends his assumption as a precisification of a vague concept, meaning that we "can understand a vague concept by looking at all the ways in which it can be made precise, and treating facts about the properties that all its "precisifications" share as facts about the concept itself."(Joyce 1998, 590). Thus, that we could understand the vague concept of accuracy by considering all precisifications of the notion of accuracy. But, as Mayo-Wilson and Wheeler write, there are many ways of making the concept of accuracy precise that do not make the concept numerical; the assumption that the precifications are numerical of Joyce is unjustified (Mayo-Wilson and Wheeler 2015, 4). Furthermore, if accuracy cannot be made numerically precise because the relation "is more accurate than" violates axioms required for a numerical representation, the precisification is not a measure of accuracy but something else. The so-called "precisification" of accuracy would then not actually be a measure of *accuracy* at all. Another argument for taking accuracy to be numerical is given by Leitgeb and Pettigrew. They write that because degrees of belief and truth-values are on the same numerical scale, accuracy or "closeness [...] to the truth" of a degree of belief can be measured "according to a metric on the one-dimensional Euclidean space"(Leitgeb and Pettigrew 2010a, 212). Supposedly, because degrees of belief are all numerical and truth values are also numerical, accuracy is a measure on the same numerical scale, where two identical values of accuracy mean that two degrees of belief are just as accurate. But just because two quantities are both numerical does not mean that they can be numerically compared: weight of an object and the volume of sound are both numerical, but there is no clear way of comparison (Mayo-Wilson and Wheeler 2015, 3). Just because my degrees of belief in two propositions—my belief of 0.34 that someone in Vinokourov's cycling team Astana is using doping and my degree of belief of 0.7 that strawberries in June (in the supermarket) are good—are both numerical does not mean, just as such, that the beliefs are just as accurate or even comparable on a single scale of accuracy. It is not clear that there is one single scale of accuracy. Of course, it is true in some sense that some degrees of belief are more accurate than other degrees of belief, but so is one Whiskey 'darker' than another in flavour—we surely would not want to say that it is 0.63 percent darker. A numerical representation of this relation, and also of accuracy, is not so easily obtained.

One reply to this might be to say that numerical epistemic utility is not needed; we are only interested in seeing when a credence function is dominated so as to make a case for the coherence conditions on degrees of belief. However, the dominance criteria is defined such that the *value* given by a scoring rule of one credence function is higher than the value of another, whatever happens. Dominance by itself is not numerical but an ordinal relation—all it says it that something dominates something else—but the definition of dominance used by accuracy-first epistemologists is based on scoring rules; degrees of belief (credence functions) are said to be dominated if there are other degrees of belief for which the *score* is uniformly less. The score therefore needs to be numerical, because otherwise the notion of dominance used by the accuracy-first epistemologists is meaningless.

A definition of dominance not based on numerical measures of accuracy (and therefore not on scores) could work here, but then the accuracy-first epistemologist would have to discard using scoring rules as measures of inaccuracy. One of the most important task for the accuracy-first theorists is to characterize dominance, because dominated degrees of belief are irrational and incoherent; only undominated degrees of belief are (represented by) probability functions. If numerical accuracy measures are discarded, the entire project has to be redone.

For this numerical representation of accuracy, which would be a numerical measure of epistemic utility, certain axioms must be satisfied by a qualitative relation[4]; here a relation between belief states at a world. Thus, using the notation introduced above, we take as primitive $(\mathbf{b}, w) \succcurlyeq (\mathbf{b}', w)$ to say that degree of belief b at world w is at least as accurate as b' is with respect to that same world. The first requirement for a numerical representation is totality: for any two beliefs $\mathbf{b}$ and $\mathbf{c}$, either $(\mathbf{b}, w) \succcurlyeq (\mathbf{c}, w)$ or $(\mathbf{c}, w) \succcurlyeq \mathbf{b}, w)$ for any world w. Whereas in traditional utility theory totality can be enforced by a forced choice measurement scenario, this is not a possible step for epistemic utility theory and hence totality has to be defended.

However, let us suppose for the sake of argument that accuracy is, as the accuracy-first epistemologists assume, numerical: how then do we then compare all these numbers? As said above, just because two measures are numerical does not mean that they are straightforwardly comparable. To take the distance from two unrelated beliefs from the truth is ambiguous. My (degree of) belief about the usage of doping in some cycling team is not clearly and unambiguously comparable to my belief about the weather. The way the different distances have to be added seems to depend on the purpose of our measurement, and this purpose of measurement is not determined by purely epistemic virtues (attaining truth, avoiding falsity, being justified in your beliefs) but by other factors as well. Accuracy, then, depends on factors other than purely epistemic matters and so fails as the foundation for Epistemic Puritanism. A similar argument goes for summing accuracy of different degrees of belief: here too, the purpose of measurement affects how we should weigh the

4. For an accessible explanation of utility theory and the axioms of Von Neumann and Morgenstern, see (Peterson 2009, 226–247).

accuracy of the different beliefs. Thus, it is not clear that accuracy is a single purely epistemic measure; it requires an argument. Perhaps it is possible to expand the list of purely epistemic virtues which are used to grade accuracy, but it is not clear that these different virtues would yield a single measure of accuracy applicable in all cases.

Transitivity poses similar problems of ambiguity in measuring accuracy. If accuracy is simply closeness to the truth, transitivity of accuracy in comparing three different agents fails. The problem is that there must be a way of *weighing* importance of degrees of belief and propositions so as to avoid intransitivity, but it is not clear that there are criteria for weighing the importance of accuracies of different degrees of belief. If such a weighing criteria can be made, it seems very likely that it has to be argued for by reference to not purely epistemic values. Here, too, there might be other purely epistemic values that could tell us how to weigh every proposition, but so far these have not been provided for by the accuracy-first epistemic puritanists.

If it is postulated that all propositions are to be weighed equally, a third requirement is violated: Independence. Independence is an axiom that says $A \succ B$ if and only if $ApC \succ BpC$, where ApC means "A combined with an event C that has a probability of p of occurring" (Peterson 2009, 99). In effect, independence says that if A is preferred to B, then it remains so if combined with some other event that has some positive probability of occuring. This is usually explained in terms of a "lottery", where a preference (say, apples over oranges) must remain so if multiplied with some chance lottery. The problem is that epistemic utility violates this axiom, too: combining degrees of belief with something like epistemic lotteries can violate the (epistemic) preference relation that one ends up with. That is, *if* there is an epistemic analogue of a lottery. Here lies another problem: defining an epistemic lottery which suffices for the independence axiom.[5] Continuity, the last axiom of the Von Neumann and Morgenstern result, is also not clearly satisfied: this would require that the there is no belief state that is infinitely better than another. But it is not clear that epistemic utility is bounded in the required way.

Summing the argument up, then, fundamental axioms are violated by an intuitive understanding of accuracy. To make for a numerical utility function, considerations of a *non-epistemic* nature seem to be needed to obtain a single numerical measure of accuracy; considerations that the Bayesian Epistemologists is exactly trying to get rid off.

5. Mayo-Wilson and Wheeler discuss as an option taking an intermediate belief state as such an epistemic lottery, which fails to satisfy independence. Furthermore, such an intermediate belief state assumes *probabilistic* degrees of belief, where probabilism is exactly what epistemic utility is meant to establish.

In a short note, Joyce has also said that his construal provides normative appeal for Scott's axiom which we have discussed in section 1.2. He writes: "Scott's axiom is just the requirement one would impose if one wanted partial beliefs to be gradationally accurate. [...] Thus, once we start thinking in terms of gradational accuracy, Scott's axiom can be interpreted as a constraint that prevents people from having partial beliefs that are less accurate than they need to be. This, as we have seen, is something to be avoided on pain of *epistemic* irrationality"(Italics in original, Joyce 1998, 602). This does not work for multiple reasons. Firstly, the uncomfortable nature of Scott's axiom lies in that it is put in non-Boolean terms, namely addition rather than union; this is a difficult condition to make sense of in considerations of likelihood. If degrees of belief are already assumed to be numerical, as Joyce does, then Scott's axiom is satisfied easily; if addition holds on degrees of belief, Scott's axiom is trivial. But if numerical degrees of belief are assumed, the awkward nature of Scott's axiom disappears, which was that it is an algebraic condition. Addition of numerical degrees of belief is also an algebraic condition. And Joyce does assume this, for the justification he gives for the condition is that it must be gradationally accurate; remember that the justification for addition was that this is necessary for maximizing gradational accuracy. Secondly, we here again require the notion of accuracy which is a concept which is not adequately defined—it is not clear that this is numerical and can work as the basis for justification of the properties of $\mathbf{P}$ or for Scott's axiom.

4.2.2 Epistemic Decision Theory

Another possible way of using scoring rules in an epistemic way that is worth considering is to obtain forecasts through an agent's decisions, and then evaluate these decision according to purely epistemic standards. This, because it would (presumably) also be constructed through usage of a measure of accuracy, would not help for the problems just discussed, but this would avoid the problem of saying what degrees of belief are, as was the problem we have discussed in the previous chapter. This theory is called *Epistemic decision theory*, as Levinstein and Konek have called their theory and how Pettigrew has also explained his accuracy-first epistemology, as applying aspects of the measure of epistemic utility to decision theory. Here, then, decisions are evaluated on the basis of their epistemic merit. This interpretation is slightly different from the approach above, where we evaluated beliefs, not acts.

This epistemic decision theory faces the same difficulties just discussed with defining the notion of epistemic utility. Furthermore, there are problems with the very notion of an epistemic decision theory; the notion of evaluating decisions on purely epistemic matters. Firstly, this requires taking any act or decision that is to be evaluated apart from its cause and reasons and evaluate it purely by itself; acts are good or bad regardless of the fulfilment of the ends an agent has set for himself. By itself, this is already odd: acts are evaluated in "goodness" or "badness" according to how well they fulfil the goals and intentions behind them. To evaluate an act according to purely epistemic standards is only applicable if the act itself is done with purely epistemic intentions or goals—if such an act is at all possible. Secondly, the idea of evaluating *actions* through their *epistemic* merit seems to make a category mistake: *beliefs* belong to epistemology and are evaluated by epistemic standards, not acts.

Here is why. Supposing that accuracy can be numerically quantified, the question is this: how do you evaluate a decision on the basis of accuracy with respect to truth? *Truth* does not apply to acts, it applies to propositions; accuracy applies to beliefs and is supposed to be a measure of the distance to the true propositions, where beliefs are propositional attitudes. Applying truth, accuracy, or any epistemic notions to decisions or acts is applying properties (if truth is not a property, being-true or being accurate surely is) to concepts where it does not even make sense to talk of the properties. Compare applying the property of being painful to the firing of neurons, or one of Ryle's original examples where a visitor, after having seen every building of the Oxford Campus, asks where the university is (Ryle 2009, 5–7).

Levinstein and Konek have not made this simple category mistake, and argue that evaluating decisions on the basis of their epistemic worth is evaluating the beliefs behind those decisions. It is not the acts that matter, but the degrees of belief that caused the act. In a sense, I have attacked a straw man, but this is straw man is worth attacking because evaluating the *acts* would avoid the problems we have seen in the previous chapter for accuracy-first epistemologists, in defining degrees of belief. The accuracy-first epistemologists usage of decision theory also makes it a reasonable interpretation.

As we have seen in this chapter, even if a clear purely epistemic explanation of the notion of a degree of belief could be given, the accuracy-first epistemologist still has to provide a clear notion of epistemic utility based on the notion of accuracy. And such a purely epistemic characterization of accuracy will not be an easy notion to define or characterize from the purely epistemic perspective.

Chapter 5

Conclusion

The problems for the synchronic part of Epistemic Puritanism are fundamental: the theory has a hard time both in defining a strong enough notion of degree of belief and in arguing for plausible coherence conditions on these degrees of belief which are required for a probability representation. If it starts with only considerations of likelihood—a "weak" notion of degree of belief—very strong conditions are required to obtain a probability representation; conditions which are all the more difficult to justify from the purely epistemic perspective. If it wants on the other hand to use a "strong" notion, where degrees of belief are numerical and every agent believes propositions to numerical degrees, it lacks a way of non-ambiguously and meaningfully ascribing numerical degrees to beliefs; it lacks a measurement scale which ascribes numbers. It is not clear what it means to have purely epistemic numerical degrees of belief. This is completely unlike the traditional account of de Finetti and Ramsey, where degrees of belief can be understood as dispositions to bet or to forecast. This account of de Finetti and Ramsey requires values, though, and is exactly because of this unavailable to the Epistemic Puritanist. Resorting to imprecise probabilities and interval-valued degrees of belief is also of no use to the Epistemic Puritanist: a measurement scale which assigns meaning to numbers, and assigns numbers in a nonambiguous way, is required both for single-valued and interval-valued degrees of belief. The strong-realist view about degrees of belief that many epistemic puritanists seem to commit to is an untenable view from precisely this purely epistemic perspective.

Furthermore, the relatively straightforward coherence conditions on numerical degrees of belief, as de Finetti introduced the conditions on betting prices and forecasts, lack a strong normative case from the purely epistemic perspective. It lacks the essential elements of De Finetti and Ramsey's account of subjective probability: values, preferences and utilities are required both in the setup and in arguing for the coherence conditions. The Bayesian epistemologist has a very hard time filling in the gaps which removing these elements leaves behind. A hopeful proposal, to substitute subjective utility with epistemic utility, also does not get of the ground: it is not clear that the notion of epistemic utility can be construed in the way some philosophers in the debate have tried.

The results of this thesis should be cause for alarm for all purely epistemic Bayesian epistemologists. Their theory is ill-founded; a notion of degree of belief is assumed which is not explicated. With the current popularity of accuracy-first epistemology, the results of this thesis are very relevant to the current philosophical debate: accuracy-first epistemology does not seem to have its foundations right. These foundations are both the notion of degree of belief and the notion of epistemic accuracy; we have seen that both these notions are fundamentally unclear from their purely epistemic perspective.

It seems to me, though, that the demand for purely epistemic subjective probability theory is also unmotivated. Firstly, the definition of subjective probability by means of scoring rules avoids many of the problems the betting argument faces, and avoids many of the criticisms philosophers have thrown at the betting interpretation and Dutch Book argument. Scoring rules already evaluate only an agent's opinions; the presence of values in the measurement and the penalty are unproblematic. It is only from the viewpoint of some forms of traditional epistemology that it seems undesirable. The relation between beliefs and actions in folk psychology is not strict, but there is a strong relation, and if one loosens the demands for purity the road to subjective probability is open.

Additionally, some other contemporary theories in "traditional" epistemology view knowledge as fundamentally a pragmatic matter, too (Fantl and McGrath 2007; Brown 2008). According to a new theory in traditional epistemology—traditional in its holding on to full beliefs and being concerned with characterizing knowledge in terms of these—knowledge is not just a matter of a belief's truth, the evidence for the belief, reliability of how the belief came to be formed, etc. (the latter two can be conceived as justification conditions), but also on the *stakes* of the belief in question. If the stakes are high, a belief does not count as knowledge very quickly; if nothing hinges very much on the belief, the belief counts as knowledge much more quickly. This theory sees knowledge as fundamentally a pragmatic matter: what matters is whether an agent should *act* on a belief; whether it is rational to act in accordance with that belief.

This theory has strong similarities to subjective probability theory as a pragmatic epistemology. *Acts* take center stage, not beliefs, and the "purely epistemic" is rejected; not only traditional "epistemic" factors matter for a belief to count as knowledge, but the pragmatic importance of the truth of that belief in whatever situation we are interested in. What exactly the similarities and differences are between the subjective probability theory here defended and this pragmatic conception of knowledge is a topic for further research. However, this theory seems to provide extra argumentation for rejecting puritanism in Bayesian Epistemology, too: not only is Epistemic Puritanist Bayesianism very problematic, it is ill-founded. If this pragmatic epistemology is correct, we should not want to try to characterize degrees of belief as rational based on "purely" epistemic criteria, but should do so on the basis of maximizing expected gain and minimizing expected loss. We should not want this puritanism, because epistemology itself is much more pragmatic than it has seemed: in fact, the puritanist project is based on a conceptual misunderstanding of "epistemology", if this pragmatic epistemology is right.

Bayesian Epistemology *with* values is also much more connected with other scientific fields outside philosophy. Subjective probability theory and decision theory have been very influential in, for example, economics and statistics, and I take this to be more than what traditional "pure" epistemology has got going for it: sceptical thought-experiments where all knowledge is doubted and rejected have very limited scientific import. This is another reason not to try going puritanist, but to keep subjective values in the picture.

These are all subjects for further discussion. I have argued in this work that a purely epistemic approach to subjective probability is, as it stands, an implausible approach. Purely epistemic subjective probabilities do not exist; but this is not a problem.

Bibliography

Anscombe, Francis John, and Robert John Aumann. 1963. “A Definition of Subjective Probability.” *The Annals of Mathematical Statistics* 34:199–205.

Anscombe, Gertrude Elizabeth Margaret. 1957. *Intention.* Harvard University Press.

Brown, Jessica. 2008. “Knowledge and Practical Reason.” *Philosophy Compass* 3:1135–1152.

Burge, Tyler. 1986. “Individualism and Psychology.” *The Philosophical Review* 95:3–45.

Carroll, Lewis. 1895. “What the Tortoise said to Achilles.” *Mind* 4:278–280.

Christensen, David. 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief.* 106–142. New York: Oxford University Press.

Churchland, Paul M. 2010. “Eliminative Materialism and the Propositional Attitudes.” In *Mind and Cognition: an Anthology,* edited by William G. Lycan and Jesse J. Prinz, 231–244. Oxford: Blackwell.

Davidson, Donald. 1980a. *Essays on Actions and Events.* Oxford: Clarendon Press.

———. 1980b. “Mental Events.” In Davidson 1980a, 207–224.

———. 1980c. “Psychology as Philosophy.” In Davidson 1980a, 229–238.

———. 1987. “Knowing One’s Own Mind.” *Proceedings and Addresses of the American Philosophical Association* 60:441–458.

De Finetti, Bruno. 1972. *Probability, Induction and Statistics: the art of guessing.* London: John Wiley & Sons.

De Finetti, Bruno. 1974. *Theory of probability: a critical introductory treatment.* Edited by Antonio Machi and Adrian Smith. Wiley.

———. 1975. *Theory of probability: a critical introductory treatment.* Edited by Antonio Machi and Adrian Smith. Wiley.

———. 1980. “Foresight: its Logical Laws, its Subjective Sources.” Edited by Henry E. Kyburg Jr. and Howard E. Smokler. Translated by Henry E. Kyburg Jr. (New York):53–118.

———. 1981. “The Role of ‘Dutch Books’ and of ‘Proper Scoring Rules’.” *The British Journal for the Philosophy of Science* 32:55–56.

———. 1993. “On the Subjective Meaning of Probability.” In *Probabilità e Induzione,* translated by Mara Khale, 291–321. Bologna: Editrice Clueb Bologna.

———. 2008. *Philosophical Lectures on Probability.* Edited by Alberto Mura. Translated by Hykel Hosni. Springer.

DeGroot, Morris H. 1970. *Optimal Statistical Decisions.* McGraw-Hill Book Company.

Dennett, Daniel C. 2002. “True Believers: The Intentional Strategy and why it Works.” In *Philosophy of Mind: Classical and Contemporary Readings,* edited by David J. Chalmers, 556–567. Oxford University Press.

deVries, Willem. 2015. “Wilfrid Sellars.” In *The Stanford Encyclopedia of Philosophy,* Summer 2015, edited by Edward N. Zalta.

Easwaran, Kenny. 2011. “Bayesianism I: Introduction and Arguments in Favor.” *Philosophy Compass:*312–320.

Easwaran, Kenny, and Brandon Fitelson. 2013. “Accuracy, Coherence and Evidence.” Draft.

Eriksson, Lina, and Alan Hájek. 2007. “What are Degrees of Belief?” *Studia Logica* 86:183–213.

Fantl, Jeremy, and Matthew McGrath. 2007. “On Pragmatic Encroachment in Philosophy.” *Philosophy and Phenomenological Research* 75:558–589.

Fishburne, Peter C. 1981. “Subjective Expected Utility: A Review of Normative Theories.” *Theory and Decision* 13:139–99.

Fishburne, Peter C. 1986. "The Axioms of Subjective Probability." *Statistical Science* 1:335–58.

Greaves, Hilary, and David Wallace. 2006. "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." *Mind* 115:607–632.

Hájek, Alan. 2008. "Arguments for—or against—Probabilism?" *The British Journal for the Philosophy of Science* 59:793–819.

Howson, Colin. 2008. "De Finetti, Countable Additivity, Consistency and Coherence." *British Journal of the Philosophy of Science* 59:1–23.

———. 2009. "Epistemic Probability and Coherent Degrees of Belief." In Huber and Schmidt-Petri 2009, 97–119.

Huber, Franz, and Christoph Schmidt-Petri, eds. 2009. Studies in Epistemology, Logic, Methodology, and Philosophy of Science 342. Springer.

Jeffrey, Richard C. 1970. "Dracula meets Wolfma: Acceptance vs. Partial Belief." In *Induction, Acceptance, and Rational Belief,* edited by Marshall Swain, 157–185. Dordrecht: D. Reidel Publishing Company.

Joyce, James M. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65:575–603.

———. 2009. "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief." In Huber and Schmidt-Petri 2009, 263–297.

Kolodny, Niko. 2007. "How does Coherence Matter?" *Proceedings of the Aristotelian Society* 107:229–263.

Kraft, Charles H., John W. Pratt, and A. Seidenberg. 1959. "Intuitive Probability on Finite Sets." *The Annals of Mathematical Statistics* 30:408–19.

Krantz, David H., Robert Duncan Luce, Patrick Suppes, and Amos Tversky. 1971. *Foundations of Measurement: Additive and Polynomial Representations.* New York: Academic Press.

Kyburg, Henry E., Jr. 1978. "Subjective Probability: Criticisms, Reflections, and Problems." *Journal of Philosophical Logic* 7:157–80.

———. 2003. "Are there degrees of belief?" *Journal of applied logic* 1:139–149.

Kyburg, Henry E., Jr., and Howard E. Smokler, eds. 1980. *Studies in Subjective Probability.* New York: Robert E. Krieger Publishing Company.

Leitgeb, Hannes, and Richard Pettigrew. 2010a. "An Objective Justification of Bayesianism 1: Measuring Inaccuracy." *Philosophy of Science* 77:201–235.

———. 2010b. "An Objective Justification of Bayesianism 2: The Consequences of Minimizing Inaccuracy." *Philosophy of Science* 77:236–272.

Levi, Isaac, and Sidney Morgenbesser. 1964. "Belief and Disposition." *American Philosophical Quarterly* 1:221–232.

Levinstein, Ben, and Jason Konek. 2015. "The Foundations of Epistemic Decision Theory." Unpublished Manuscript.

MacFarlane, John. 2004. "In What Sense (if any) is Logic Normative for Thought?" *Delivered at the American Philosophical Association Central Division Meeting.*

Maher, Patrick. 2002. "Joyce's Argument for Probabilism." *Philosophy of Science* 69:73–81.

Mayo-Wilson, Conor, and Gregory Wheeler. 2015. "Epistemic Decision Theory's Reckoning." Unpublished Manuscript; under review.

Meacham, Christopher J. G., and Jonathan Weisberg. 2011. "Representation theorems and the Foundations of Decision Theory." *Australasian Journal of Philosophy* 89:641–663.

Millikan, Ruth Garrett. 1993. *White Queen Psychology and Other Essays for Alice.* Cambridge: MIT Press.

Peterson, Martin. 2009. *An Introduction to Decision Theory.* Cambridge: Cambridge University Press.

Pettigrew, Richard. 2013. "Epistemic Utility and Norms for Credences." *Philosophy Compass* 8:897–908.

Ramsey, Frank Plumpton. 1926. "Truth and Probability." In Kyburg and Smokler 1980, 23–52.

Ryle, Gilbert. 2009. *The Concept of Mind.* New York: Routledge.

Savage, Leonard J. 1954. *The Foundations of Statistics.* Wiley.

Schwitzgebel, Eric. 2015. “Belief.” In *The Stanford Encyclopedia of Philosophy,* Summer 2015, edited by Edward N. Zalta.

Scott, Dana. 1964. “Measurement Structures and Linear Inequalities.” *Journal of Mathematical Psychology* 1:233–47.

Seidenveld, Teddy. 2014. *Dominance and Elicitation in IP Theory.* Itunes U presentation: Munich Center of Mathematical Philosophy.

Sellars, Wilfrid. 1956. “Empiricism and the Philosophy of Mind.” In *The Foundations of Science and the Concepts of Psychology and Psychoanalysis,* vol. 1 of *Minnesota Studies in the Philosophy of Science,* edited by Herbert Feigl and Michael Scriven, 253–329. University of Minnesota Press.

Suppes, Patrick. 2002. *Representation and Invariance of Scientific Structures.* Stanford: CSLI Publications.

Suppes, Patrick, and Mario Zanotti. 1976. “Necessary and Sufficient Conditions for Existence of a Unique Measure Strictly Agreeing with a Qualitative Probability Ordering.” *Journal of Philosophical Logic* 5:431–438.

Vicig, Paolo, and Teddy Seidenfeld. 2012. “Bruno De Finetti and Imprecision: Imprecise Probability does not exist!” *International Journal of Approximate Reasoning* 53:1115–1123.

Walley, Peter. 1991. *Statistical Reasoning with Imprecise Probabilities.* London: Chapman & Hall.

Williamson, Jon. 1999. “Countable Additivity and Subjective Probability.” *British Journal of the Philosophy of Science* 50:401–16.

Zynda, Lyle. 2000. “Representation Theorems and Realism about Degrees of Belief.” *Philosophy of Science* 67:45–69.