Belief and Degrees of Belief

Franz Huber, California Institute of Technology

Draft June 12, 2006

Contents

In memory of Philippe Smets (1938-2005)

1 Introduction

Degrees of belief are familiar to all of us. Our confidence in the truth of some propositions is higher than our confidence in the truth of other propositions. We are pretty confident that our computers will boot when we push their power button, but we are much more confident that the sun will rise tomorrow. Degrees of belief formally represent the strength with which we believe in the truth of various propositions. The higher an agent's degree of belief for a particular proposition, the higher her confidence in the truth of the proposition. For instance, Sophia's degree of belief that it will be sunny in Vienna tomorrow might be .69, whereas her degree of belief that the train will leave on time might be .23. The precise meaning of these statements depends, of course, on the underlying theory of degrees of belief. These theories offer a formal tool to measure degrees of belief, to investigate the relations between various degrees of belief in different propositions, and to normatively evaluate degrees of belief.

The purpose of this book is to provide a comprehensive overview and assessment of the currently prevailing theories of degrees of belief. Degrees of belief are primarily studied in epistemology. The aim is to adequately describe and, much more importantly, to normatively evaluate the epistemic state of an (ideal or rational) agent. Here theories of degrees of belief allow a precise analysis that is hard to come by with traditional philosophical methods. Degrees of belief are also studied in computer science and artificial intelligence, where they find applications in so-called expert systems and elsewhere.

Different theories of degrees of belief postulate different ways in which degrees of beliefs are related to each other and, more generally, how epistemic states should be modeled. After getting a handle on the objects of belief in section 2, we briefly survey the most important accounts in section 3. Section 4 continues this survey by focusing on the relation between belief and degrees of belief. Section 5 concludes this introduction by pointing at some relations to belief revision and nonmonotonic reasoning.

2 The Objects of Belief

Before we can investigate the relations between various degrees of belief, we have to get clear about the relata of the (degree of) belief relation. It is common to assume that belief is a relation between an epistemic agent at a particular time to an object of belief. Degree of belief is then a relation between a number, an epistemic agent at a particular time, and an object of belief. It is more difficult to state what these objects of belief are. Are they sentences or propositions expressed by sentences or possible worlds (whatever *these* are – Stalnaker 1984) or something altogether different?

The received view is that the objects of belief are propositions, i.e. sets of possible worlds or truth conditions. Most people stay very general and assume only that there is a non-empty set of possibilities W such that exactly one element of W corresponds to the actual world. The set of all possibibilities W is a proposition, and if A and B are propositions, then so are the complement of A (with respect to W), $W \setminus A = \overline{A}$, as well as the intersection of A and B, $A \cap B$. In other words, the set of propositions is a (finitary) field $\mathcal{A}$ over a non-empty set of possibilities W, i.e. a set that contains W and is closed under complementations and finite intersections. Sometimes propositions are not only assumed to be closed under finite, but under countable intersections. This means that $A_1 \cap \ldots \cap A_n \ldots$ is a proposition (an element of $\mathcal{A}$), if $A_1, \ldots, A_n \ldots$ are. Such a field $\mathcal{A}$ is called a *σ-field*.

If Sophia believes (to some degree) that it will be sunny in Vienna tomorrow, but she does not believe (to the same degree) that it will not be not sunny in Vienna tomorrow, propositions cannot be the objects of Sophia's (degrees of) belief(s). That it will be sunny in Vienna tomorrow and that it will not be not sunny in Vienna tomorrow is one and the same proposition. It is only expressed by two different, though logically equivalent sentences.

In some accounts sentences of a formal language $\mathcal{L}$ are taken to be the objects of belief. The tautological sentence τ is assumed to be in the language $\mathcal{L}$, and whenever α and β are in $\mathcal{L}$, then so is the negation of α, $\neg\alpha$, as well as the conjunction of α and β, $\alpha \wedge \beta$. However, as long as logically equivalent sentences are required to be assigned the same degree of belief – and all accounts considered in this volume do require so – the difference between taking the objects of beliefs to be sentences from a language $\mathcal{L}$ or to be propositions from a finitary field $\mathcal{A}$ is mainly cosmetic. Each language $\mathcal{L}$ induces a finitary field $\mathcal{A}$ over the set of all models or classical truth value assignments for $\mathcal{L}$, $Mod_{\mathcal{L}}$. It is simply the set of all propositions over $Mod_{\mathcal{L}}$ that are expressed by the sentences in $\mathcal{L}$. This set in turn induces a σ-field, viz. the smallest σ-field $\sigma(\mathcal{A})$ that contains $\mathcal{A}$. Hence, if we start with a degree of belief function on a language $\mathcal{L}$, we automatically get a degree of belief function on the field $\mathcal{A}$ induced by $\mathcal{L}$. As the converse is not true, the semantic framework of propositions is simply more general than the syntactic framework of sentences.

The models of a language $\mathcal{L}$ in $Mod_{\mathcal{L}}$ are also called truth conditions, because they determine whether a sentence – that is, a closed formula of a formal language – is true or false in them. According to an alternative view (Spohn 1997) the objects of belief are sets of *satisfaction* conditions rather than sets of truth conditions. Satisfaction conditions consist of a truth condition together with a list of objects, and they determine whether an open formula of a formal language is satisfied by them. However, as long as the objects of belief form a field $\mathcal{A}$ over W, the exact nature of the elements in the set of possibilities W one level below does not put any restrictions on the laws degrees of belief have to satisfy. Moreover, the stance on the objects of belief in various accounts does not so much differ in what are taken to be the elements of W. They differ primarily in whether the objects of belief are elements of W, sentences from $\mathcal{L}$, or propositions of $\mathcal{A}$.

3 Theories of Degrees of Belief

We have started with the example of Sophia, whose degree of belief that it will be sunny in Vienna tomorrow equals .69. Usually degrees of belief are taken to be real numbers from the interval $[0, 1]$ (but we will come across an alternative in section 4). If and only if the epistemic agent is certain that a proposition is true, her degree of belief for this proposition is 1. If the epistemic agent is certain that a proposition is false, her degree of belief for the proposition is 0. The converse is true for subjective probabilities, but not in general. For instance, Dempster-Shafer belief functions behave differently, because there one distinguishes between a belief function and a plausibility function (for more see below). However, these are extreme cases. Usually we are neither certain that a proposition is true nor that it is false. That does not mean, though, that we are agnostic with respect to the question whether the proposition in question is true. Our belief that it is true may well be much stronger than that it is false. Degrees of belief are supposed to quantify this strength.

3.1 Subjective Probabilities

By far the best developed account of degrees of belief is the theory of subjective probabilities. Degrees of belief would, on this view, simply follow the laws of probability. For this reason it is also called *probabilism*. Here is the standard definition due to Kolmogorov (1956). Let $\mathcal{A}$ be a field of propositions over the set of possibilities W. A function $\Pr : \mathcal{A} \rightarrow \Re$ from $\mathcal{A}$ into the set of real numbers $\Re$ is a (*finitely additive* and *unconditional*) *probability* on $\mathcal{A}$ if and only if for all A, B in $\mathcal{A}$:

1. $\Pr(A) \geq 0$
2. $\Pr(W) = 1$
3. $\Pr(A \cup B) = \Pr(A) + \Pr(B)$ if $A \cap B = \emptyset$

The triple $\langle W, \mathcal{A}, \Pr \rangle$ is called a (*finitely additive*) *probability space*. If $\mathcal{A}$ is closed under countable intersections and thus a σ-field, and if $\Pr$ additionally satisfies

4. $\Pr(A_1 \cup \ldots \cup A_n \cup \ldots) = \Pr(A_1) + \ldots + \Pr(A_n) + \ldots$

$\Pr$ is a σ- or *countably additive probability* on $\mathcal{A}$ (Kolmogorov 1956, ch. 2, actually gives a different but equivalent definition – see e.g. Huber 2006b, sct. 4.1). In this case $\langle W, \mathcal{A}, \Pr \rangle$ is called a σ- or *countably additive probability space*.

A probability $\Pr : \mathcal{A} \rightarrow \Re$ on $\mathcal{A}$ is called *regular* just in case $\Pr(A) > 0$ for every non-empty A in $\mathcal{A}$. Let $\mathcal{A}^*$ be the set of all propositions A from $\mathcal{A}$ with $\Pr(A) > 0$. The *conditional probability* $\Pr(\cdot \mid \cdot) : \mathcal{A} \times \mathcal{A}^* \rightarrow \Re$ on $\mathcal{A}$ (based on the unconditional probability $\Pr : \mathcal{A} \rightarrow \Re$ on $\mathcal{A}$) is defined for all A in $\mathcal{A}$ and all $B \in \mathcal{A}^*$ by the fraction

5. $\Pr(A \mid B) = \Pr(A \cap B) / \Pr(B)$

(Kolmogorov 1956, ch. 1, §4). The domain of the second argument place of $\Pr(\cdot \mid \cdot)$ has to be restricted to $\mathcal{A}^*$, since the fraction $\Pr(A \cap B) / \Pr(B)$ is not defined for $\Pr(B) = 0$. Note that $\Pr(\cdot \mid B) : \mathcal{A} \rightarrow \Re$ is a probability on $\mathcal{A}$, for every $B \in \mathcal{A}^*$. Other authors take conditional probability as primitive and define unconditional probability in terms of it (Hájek 2003).

In terms of subjective probabilities Sophia's having a degree of .69 for the proposition that it will be sunny in Vienna tomorrow means that she considers the following bet to be fair: Sophia gets 31 Cents from the Bookie if it is sunny; otherwise she pays him 69 Cents. To consider this bet to be fair means that Sophia is indifferent as to which side of the bet to take, i.e. whether to bet on or against the proposition that it will be sunny in Vienna tomorrow. It is presupposed that Sophia is not risk averse, and that the monetary gains are proportional to her utilities. In real life these idealizing assumptions are not met. For gamblers risk often has a positive utility, and for many people it generally has a negative utility. Moreover, 31 Cents are a lot for somebody who lacks 30 cents for a bus ride home, whereas it means hardly anything to a wealthy man.

The statement that Sophia's having a subjective probability of .69 *means* that she considers a particular bet to be fair requires some qualifications. Sometimes subjective probabilities are *defined* as fair betting ratios. Sometimes they are merely *measured* by fair betting ratios. In this latter case we face the question how exactly subjective probabilities and fair betting ratios are related to each other. This will become important below.

The theory of subjective probabilities does not (primarily) aim at an adequate description of people's epistemic states. First and foremost it is a normative account that tells us how an ideal or rational epistemic agent's degrees of belief should behave. So, why should such an agent's degrees of belief obey the probability calculus?

The *Dutch Book Argument* (hereafter DBA) provides an answer to this question. (Cox's theorem, Cox 1961, and the representation theorem of measure theory, Krantz & Luce & Suppes & Tversky 1971, provide two further answers. See Howson's contribution.) The DBA starts with the assumption that there is an intimate link between subjective degrees of belief and fair betting ratios. It is further taken for granted that it is (pragmatically) defective to accept a bet which guarantees a sure loss. The core of the argument is the *Dutch Book Theorem* according to which an agent's fair betting ratios satisfy the probability axioms just in case they do not make the agent vulnerable to a sure loss. From this it is inferred that it is (epistemically) defective to have degrees of belief that violate the probability axioms. The strength of this argument is, of course, dependent on the link between degrees of belief and fair betting ratios. If this link is identity – as it is when one defines degrees of belief as fair betting ratios – the distinction between pragmatic and epistemic defectiveness disappears, and the DBA is a deductively valid argument, provided the Dutch Book Theorem is (Hájek 2005). But this comes at the cost of rendering the link between degrees of belief and fair betting ratios less plausible. If the link is weaker than identity – as it is when degrees of belief are only measured by fair betting ratios – the DBA is not deductively valid, but it has more plausible premisses. Joyce (1998) provides a different justification of probabilism in terms of the accuracy of degrees of belief (see his contribution). For more on the DBA, Joyce's non-pragmatic vindication of probabilism, and arguments for (non-) probabilism in general see Hájek's contribution.

We have seen how subjective probabilities have to be related to each other as well as how to measure them, and why degrees of belief should obey the probability calculus. It is also important to know how to update one's subjective probabilities when new information is received. Whereas axioms 1-5 of the probability calculus are *synchronic* conditions on an agent's degree of belief function, update rules are *diachronic* conditions that tell us how to revise our subjective probabilities when we receive new information of a certain format. If the new information comes in form of a proposition that one strictly speaking learns, probabilism is extended by

Update Rule 1 (Strict Conditionalization) *If* $\Pr : \mathcal{A} \rightarrow \Re$ *is your subjective probability at time* t*, and between* t *and* t' *you learn* $E \in \mathcal{A}$ *and no logically stronger proposition, then your subjective probability at time* t' *should be* $\Pr(\cdot \mid E) : \mathcal{A} \rightarrow \Re$.

So strict conditionalization says that the agent's new subjective probability for a proposition A after strictly learning E should be equal to her old subjective probability for A conditional on E.

Two questions arise. First, why should we update our subjective probabilities according to strict conditionalization? In particular, can we justify strict conditionalization along the lines of the DBA? Second, how should we update our subjective probabilities when the new information is of a different format and we do not strictly speaking learn a proposition, but change our subjective probabilities for various propositions? The second question is answered by

Update Rule 2 (Jeffrey Conditionalization) *If* $\Pr : \mathcal{A} \rightarrow \Re$ *is your subjective probability at time* t*, and between* t *and* t' *your subjective probabilities in the mutually exclusive and jointly exhaustive propositions* $E_1, \ldots, E_n, \ldots$ *(*$E_i \in \mathcal{A}$*) change to* $p_i \in [0, 1]$ *with* $\sum_i p_i = 1$*, and the positive part of your subjective probability does not change on any superset thereof, then your subjective probability at time* t' *should be* $\Pr' : \mathcal{A} \rightarrow \Re$*, where*

$$\Pr{}'(\cdot) = \sum\nolimits_i \Pr(\cdot \mid E_i) \cdot p_i.$$

So Jeffrey conditionalization says that the agent's new subjective probability for A after changing her subjective probabilities for the elements of a partition E_i to p_i should be equal to the weighted sum of her old subjective probabilities for A

conditional on the E_i, where the weights are the new subjective probabilities for the E_i. For the first question as well as its extension to Jeffrey conditionalization we refer to Skyrms' contribution (see also Skyrms 1987).

In subjective probability theory complete ignorance of the epistemic agent with respect to a particular proposition A is often modeled by the agent's having a subjective probability of .5 for A as well as its complement $\overline{A}$; or, more generally, as having a subjective probability of $1/n$ for each of the mutually exclusive and jointly exhaustive propositions $A_1, \ldots, A_n$ about a particular subject matter. This is known as the *principle of indifference*. It leads to contradictory results if the partition in question is not held fixed.

Furthermore, suppose Sophia has hardly any enological knowledge. Then her subjective probability for the proposition that a Schilcher is a white wine might reasonably be .5 – as might be her subjective probability that a Schilcher is a red wine. Contrast this with the following case. Suppose Sophia knows that a particular coin is fair – that is, Sophia knows that the objective chance of heads as well as tails is .5. Then her subjective probability that the coin, if tossed, will land heads might equally reasonably be .5. Although Sophia's subjective probabilities are alike in these two scenarios, there is an important difference. In the one case a subjective probability of .5 represents ignorance, in the other case it represents substantial knowledge about the objective chances. (The principle that one's subjective probabilities conditional on the objective chances should equal the objective chances is called the *Principal Principle* by Lewis 1980.)

Examples like these suggest that subjective probability theory does not provide an adequate account of degrees of belief, because it does not allow one to distinguish between ignorance and knowledge about chances. *Interval-valued probabilities* can be seen as a reply to this objection without giving up the probabilistic framework. In case one knows the objective chances one continues to assign sharp probabilities as usual. However, if the agent is ignorant with respect to a proposition A she will not assign it subjective probability .5 (or any other sharp value, for that matter). Rather, she will assign A a whole interval $[a, b] \subseteq [0, 1]$ such that she considers any number in $[a, b]$ to be a legitimate subjective probability for A. The size of the interval $[a, b]$ reflects her ignorance about A. If she assigns $[0, 1]$ to A, she is completely ignorant about A – any subjective probability is considered to be legitimate. If the interval is narrower, say $[.4, .6]$ she is still somewhat ignorant, but much less so as in the previous case. For more see Kyburg's contribution.

In science as well as in our more ordinary epistemic enterprises we constantly talk about our beliefs – those propositions which we think are true – without mentioning a particular degree of belief representing the strength of our beliefs. But what exactly is the relation between such beliefs and degrees of belief? A very natural thesis – the *Lockean thesis* – says that we (should) believe in a proposition A just in case our degree of belief for A is sufficiently high. But how high is sufficiently high? We do not want to require that we only believe those propositions which receive the highest possible degree of belief. We want to take into account our fallibilism, the fact that our beliefs often turn out to be false.

Let us assume once more that degrees of belief are represented by subjective probabilities. Then the above means that the threshold level for belief should not be 1. In particular, this should be so if we follow Carnap (1962) and Jeffrey (2004) and assign probability 1 only to tautologies – that is, if we adopt a regular probability as our degree of belief function (note, though, that this is not always possible). Otherwise we would only believe the tautology.

The Lockean thesis can therefore be spelt out in terms of subjective probabilities as follows. An epistemic agent with subjective probability $\Pr : \mathcal{A} \rightarrow \Re$ believes $A \in \mathcal{A}$ if and only if $\Pr(A) \geq 1 - \varepsilon$, for some $\varepsilon \in (0, 1]$. This, however, leads to the following *lottery paradox* (Kyburg 1961, Hempel 1962; see also Kyburg 1970 as well as the preface paradox due to Makinson 1965). Whatever threshold value $\varepsilon \in (0, 1]$ we choose, there will always be propositions $A_1, \ldots, A_n$ from some field $\mathcal{A}$ and a reasonable subjective probability $\Pr : \mathcal{A} \rightarrow \Re$ on $\mathcal{A}$ such that $\Pr(A_1) \geq 1 - \varepsilon, \ldots, \Pr(A_n) \geq 1 - \varepsilon$, and $\Pr(A_1 \cap \ldots \cap A_n) < 1 - \varepsilon$. For instance, let $\varepsilon = .02$ and consider a lottery with 100 tickets that is known to be fair. It is reasonable to have $\Pr(T_1 \cup \ldots \cup T_n) = 1$ as well as $\Pr(T_i) = .01$, where T_i is the proposition that ticket i will win, $i = 1, \ldots, 100$. This entails $\Pr\left(\overline{T_1} \cap \ldots \cap \overline{T_n}\right) = 0$ and $\Pr\left(\overline{T_i}\right) = .99$ for $i = 1, \ldots, 100$, where $\overline{T_i}$ is the proposition that ticket i will not win. Hence, according to our analysis, an epistemic agent with $\Pr$ as her subjective probability will believe that at least one of the tickets $1, \ldots, 100$ will win, and at the same time she will believe of each ticket that it will not win. This means that our agent's *belief set*, the set of all propositions she believes, is inconsistent or not deductively closed. Yet consistency and deductive closure are the minimal requirements on a belief set (Hintikka 1961).

The lottery paradox has led some people to reject the notion of belief altogether (Jeffrey 1970), whereas others have been lead to the idea that belief sets need not be deductively closed (Foley 1992). Still others have turned the analysis on its head and elicit the context dependent threshold parameter ε from the agent's belief set (Hawthorne and Bovens 1999). For more see the contributions by Foley and by Hawthorne. Another view is to take the lottery paradox at face value and postulate two epistemic attitudes towards propositions – belief and degrees of belief – that are not reducible to each other. Frankish (2004) defends a particular version of this view. He distinguishes between a mind, where one unconsciously entertains beliefs, and a supermind, where one consciously entertains beliefs. For more see his contribution.

3.2 Dempster-Shafer Belief Functions

The theory of *Dempster-Shafer* (*DS*) *belief functions* (Dempster 1968, Shafer 1976) rejects the claim that degrees of belief can be measured by the epistemic agent's betting behavior. As a consequence, they need not obey the probability calculus either.

A particular version of the theory of DS belief functions is the *transferable belief model* (Smets & Kennes 1994). It distinguishes between two mental levels: the credal and the pignistic level. Its twofold thesis is that fair betting ratios should indeed obey the probability calculus, but that degrees of belief – being different from fair betting ratios – need not. It suffices that they satisfy the weaker DS principles. So, whenever one is forced to bet on the pignistic level, degrees of belief are used to calculate fair betting ratios that satisfy the probability axioms. These are then used to calculate the agent's expected utility for various acts (Savage 1972). However, on the credal level where one only entertains and quantifies various beliefs without using them for decision making, degrees of belief need not obey the probability calculus.

Whereas subjective probabilities are additive (axiom 3), DS belief functions $Bel : \mathcal{A} \to \Re$ are only super-additive, i.e. for all $A, B \in \mathcal{A}$:

6. $Bel\,(A) + Bel\,(B) \leq Bel\,(A \cup B)$ if $A \cap B = \emptyset$.

This means in particular that the degree of belief for A and the degree of belief for $\overline{A}$ need not sum to 1.

What does it mean that Sophia's degree of belief that it will be sunny in Vienna tomorrow is .69, if her degree of belief function is represented by a DS belief function $Bel : \mathcal{A} \to \Re$? According to one interpretation (Haenni & Lehmann 2003), the number $Bel\,(A)$ represents the strength with which A is supported by the epistemic agent's knowledge base. It may well be that the epistemic agent's knowledge base neither supports A nor its complement $\overline{A}$. Recall that Sophia has hardly any enological knowledge. So her knowledge will neither support the proposition R that a Schilcher is a red wine nor will it support the proposition W that it is a white wine. Hence her DS belief function Bel will be such that $Bel\,(R) = Bel\,(W) = 0$. On the other hand, as Sophia knows the coin to be fair, her Bel will be such that $Bel(H) = Bel(T) = .5$. Thus we see that the theory of DS belief functions can easily distinguish between igorance and uncertainty. Indeed,

$$I(\mathrm{A}) = 1 - Bel(A_1) - \ldots - Bel(A_n) - \ldots$$

can be seen as a measure of the agent's ignorance with respect to the partition $\mathrm{A} = \{A_1, \ldots, A_n, \ldots\}$ (the A_i may, for instance, be the values of a random variable).

Figuratively, a proposition A divides the agent's knowledge base into three mutually exclusive and jointly exhaustive parts. A part that speaks in favor of A, a part that speaks against A (i.e. in favor of $\overline{A}$), and a part that neither speaks in favor of nor against A. $Bel(A)$ quantifies the part that supports A, $Bel\left(\overline{A}\right)$ quantifies the part that supports $\overline{A}$, and $I(\mathrm{A}) = 1 - Bel(A) - Bel\left(\overline{A}\right)$ quantifies the part that neither supports A nor $\overline{A}$. Formally this is spelt out in terms of a (normalized) *mass function* on $\mathcal{A}$, a function $m : \mathcal{A} \rightarrow \Re^+$ such that

$$m(\emptyset) = 0, \quad \sum_{A \in \mathcal{A}} m(A) = 1.$$

A (normalized) mass function $m : \mathcal{A} \rightarrow \Re^+$ induces a DS belief function $Bel : \mathcal{A} \rightarrow \Re$ by defining for each $A \in \mathcal{A}$,

$$Bel(A) = \sum_{B \subseteq A} m(B).$$

The relation to subjective probabilities can now be stated as follows. Subjective probabilities require the epistemic agent to divide her knowledge base into two mutually exclusive and jointly exhaustive parts: one that speaks in favor of A and one that speaks against A. That is, the neutral part has to be distributed among the positive and negative parts. Thus subjective probabilities can be seen as DS belief functions without ignorance.

A DS belief function $Bel : \mathcal{A} \rightarrow \Re$ induces a *Dempster-Shafer plausibility function* $P :\rightarrow \Re$, where for all $A \in \mathcal{A}$:

$$P(A) = 1 - Bel(A).$$

Degrees of plausibility quantify that part of the agent's knowledge base which is compatible with A, i.e. the part that supports A and the part that supports neither A nor $\overline{A}$. In terms of the (normalized) mass function inducing Bel this means that

$$P(A) = \sum_{B \not\subseteq \overline{A}} m(B).$$

While it is possible that $Bel\,(A)$ and $Bel\,(\overline{A})$ sum to less than 1, it may happen that $P\,(A)$ and $P\,(\overline{A})$ sum to more than 1. For more see Haenni's contribution.

Although we have just said that the theory of DS belief functions is more general than the theory of subjective probability, in another sense the converse is true. The reason is that DS belief functions can be represented as *convex sets of probabilities* (Levi 1980, Walley 1991). The same is true for interval-valued probabilities. As not every convex set of probabilities can be represented as a DS belief function (or an interval-valued probability), sets of probabilities provide the most general framework we have come across so far. A similarly general framework is provided by Halpern's plausibility measures (Halpern 2003). These are functions $Pl : \mathcal{A} \rightarrow \Re$ such that for all $A, B \in \mathcal{A}$: $Pl\,(\emptyset) = 0$, $Pl\,(W) = 1$, and

7. $Pl\,(A) \leq Pl\,(B) \quad \text{if} \quad A \subseteq B.$

In fact, this is only the special case of real-valued plausibility measures. The problem, though, is that such general models do not allow very specific claims about the agent's epistemic state. For more see Halpern's contribution.

3.3 Possibility Theory

Possibility theory (Dubois & Prade 1988) is based on fuzzy set theory (Zadeh 1978). According to the latter theory, an element need not belong to a set either completely or not at all, but may be a member of the set to a certain degree. For instance, Sophia may belong to the set of black haired women to a degree of .69, because her hair, although black, is sort of brown as well (we can assure you that she belongs to the set of beautiful women to degree 1). This is represented by a membership function $\mu_B : W \rightarrow [0, 1]$, where $\mu_B\,(w)$ is the degree of membership to which woman $w \in W$ belongs to the set of red haired woman B.

If $\mu_B\,(\text{Sophia}) = .69$, then the degree $\mu_{\overline{B}}\,(\text{Sophia})$ to which Sophia belongs to the set $\overline{B}$ of women who do not have black hair equals $1 - \mu_B\,(\text{Sophia})$. Moreover, if $\mu_Y : W \rightarrow [0, 1]$ is the membership function for the set of young women, then the degree of membership to which Sophia belongs to the set $B \cup Y$ of black haired or young women is given by

$$\mu_{B \cup Y}\,(\text{Sophia}) = \max\left\{\mu_B\,(\text{Sophia})\,, \mu_Y\,(\text{Sophia})\right\}.$$

Similarly, the degree of membership to which Sophia belongs to the set $B \cap Y$ of black haired young women is given by

$$\mu_{B \cap Y}\,(\text{Sophia}) = \min\left\{\mu_B\,(\text{Sophia})\,, \mu_Y\,(\text{Sophia})\right\}.$$

μ_B (Sophia) is interpreted as the degree to which the vague statement 'Sophia is a black haired woman' is true.

Degrees of truth belong to ontology and philosophy of language. They do not (yet) have anything to do with degrees of belief, which belong to epistemology. In particular, note that degrees of truth are usually considered to be truth functional (the truth value of a compound statement is a function of the truth values of its constituent statements). This is the case for membership functions μ. Degrees of belief, on the other hand, are hardly ever considered to be truth functional. For instance, probabilities are not truth functional, because the probability of $A \cap B$ is not determined by the probability of A and the probability of B.

Suppose we tell you that Sophia is tall. How tall is a tall woman? Is a woman with a body height of 175cm tall? Or does a woman have to have a body height of at least 178cm in order to be tall? Although you know that Sophia is tall, the vagueness of the term 'tall' prevents you from knowing how tall Sophia is. In other words, your knowledge is incomplete. Here possibility theory enters by equipping you with a (normalized) *possibility distribution*, a function $\pi : W \to [0,1]$ with $\pi(\omega) = 1$ for at least one $\omega \in W$. The motivation for the latter requirement is that at least one possibility is the actual possibility, and hence at least one possibility must be maximally possible. Such a possibility distribution $\pi : W \to [0,1]$ on the set of possibilities W is readily extended to a *possibility measure* $\Pi : \mathcal{A} \to \Re$ on the field $\mathcal{A}$ over W by defining for each $A \in \mathcal{A}$,

$$\Pi(\emptyset) = 0, \quad \Pi(A) = \sup\{\pi(\omega) : \omega \in A\}.$$

This entails that possibility measures $\Pi : \mathcal{A} \to \Re$ are *maxitive* (and hence sub-additive), i.e. for all $A, B \in \mathcal{A}$:

8. $\Pi(A \cup B) = \max\{\Pi(A), \Pi(B)\}$.

The idea is, roughly, that a proposition is at least as possible as all of the possibilities it comprises, and no more possible than the "most possible" possibility either. Sometimes, though, there is no most possible possibility (i.e. the supremum is no maximum) – e.g. when the degrees of possibility are 1/2, 3/4, 7/8, $\ldots, \frac{2^n-1}{2^n}, \ldots$ In this case the degree of possibility for the proposition is the smallest number which is at least as great as all the degrees of possibilities of its elements – in our example this is 1. (As will be seen below, this is the main formal difference between possibility measures and ranking functions.)

The dual notion of a *necessity measure* $N : \mathcal{A} \to \Re$ is defined for all $A \in \mathcal{A}$ by

$$N(A) = 1 - \Pi\left(\overline{A}\right).$$

Although the agent's epistemic state is completely specified by either Π or N, the agent's epistemic attitude towards a particular proposition $A \in \mathcal{A}$ is only jointly specified by $\Pi(A)$ and $N(A)$. The reason is that, in contrast to probability theory, $\Pi\left(\overline{A}\right)$ is not determined by $\Pi(A)$. Thus, degrees of possibility (as well as degrees of necessity) are not truth functional either. The same is true for DS belief and plausibility functions.

In our example, let W_H be the set of values of the random variable H = *Sophia's body height in cm between 0cm and 300cm*, $W_H = \{0, \ldots, 300\}$. $\pi_H : W_H \to [0, 1]$ is your possibility distribution. It is supposed to represent your epistemic state concerning Sophia's body height, which contains the knowledge that she is tall. For instance, your π_H might be such that $\pi_H(n) = 1$ for any natural number $n \in [177, 185] \subset W$. In this case your degree of possibility for the proposition that Sophia is at least 177cm tall is $\Pi_H(H \geq 177) = \sup\{\pi_H(n) : n \geq 177\} = 1$.

The connection to fuzzy set theory now is that your possibility distribution $\pi_H : W_H \to [0, 1]$, which is based on the knowledge that Sophia is tall, can be interpreted as the membership function $\mu_T : W_H \to [0, 1]$ of the set of tall woman. So the epistemological thesis of possibility theory is that your degree of possibility for propositions such as 'Sophia is 177cm tall' given (due to its vagueness) incomplete knowledge such as 'Sophia is tall' equals the degree of membership to which a 177cm tall woman belongs to the set of tall woman. In more suggestive notation,

$$\pi_H(H = n \mid T) = \mu_T(n).$$

For more see the contribution by Dubois and Prade.

3.4 Summary

One difference between the three accounts we have dealt with so far is the following. Subjective probability theory requires degrees of belief to be additive. An ideal or rational epistemic agent's subjective probability $\Pr : \mathcal{A} \to \Re$ is such that for any $A, B \in \mathcal{A}$ with $A \cap B = \emptyset$:

$$\Pr(A) + \Pr(B) = \Pr(A \cup B).$$

The theory of DS belief functions requires degrees of belief to be super-additive. An ideal or rational epistemic agent's DS belief function $Bel : \mathcal{A} \to \Re$ is such that for any $A, B \in \mathcal{A}$ with $A \cap B = \emptyset$:

$$Bel\,(A) + Bel\,(B) \leq Bel\,(A \cup B)\,.$$

Possibility theory requires degrees of belief to be maxitive. An ideal or rational epistemic agent's possibility measure $\Pi : \mathcal{A} \to \Re$ is such that for any $A, B \in \mathcal{A}$:

$$\Pi\,(A \cup B) = \max\left\{\Pi\,(A)\,, \Pi\,(B)\right\}.$$

All of these functions are special cases of real-valued plausibility measures Pl : $\mathcal{A} \to \Re$, which are such that for all $A, B \in \mathcal{A}$:

$$Pl\,(A) \leq Pl\,(B) \quad \text{if} \quad A \subseteq B.$$

We have seen that each of these accounts provides an adequate model for some epistemic situation (plausibility measures do so trivially). We have further noticed that subjective probabilities do not give rise to a notion of belief that is consistent and deductively closed. The same is, of course, true for the more general DS belief functions and plausibility measures. It has to be noted, though, that Roorda (1995) provides a definition of belief in terms of sets of probabilities. As will be mentioned in the next section, the situation is different for possibility theory.

Moreover, we have seen arguments for the thesis that degrees of belief should obey the probability axioms. Smets (2002) tries to justify the corresponding thesis for DS belief functions. To the best of our knowledge nobody has yet published an argument for the thesis that degrees of belief should be plausibility measures (not just in the sense that only plausibility measures are reasonable degree of belief functions, but in the sense that *all* and only plausibility measures are reasonable degree of belief functions.) We are not aware of an argument for possibility measures as degree of belief functions either. However, there exists such an argument for the formally similar ranking functions. The latter functions also give rise to a notion of belief that is consistent and deductively closed (indeed, this very feature is the basis for the argument that epistemic states should be modeled by ranking functions). Ranking functions are the topic of the next section.

4 Belief, Degrees of Belief, and Ranking Functions

Subjective probability theory as well as the theory of DS belief functions take the objects of belief to be propositions. Possibility theory does so only indirectly, though possibility measures on a field of propositions $\mathcal{A}$ can also be defined without recourse to a possibility distribution on the underlying set of possibilities W.

A possibility $\omega \in W$ is a complete and consistent description of what the world looks like relative to the expressive power of W. W may contain two possibilities: According to ω_1 it will be sunny in Vienna tomorrow, according to ω_2 it will not. Or else, W may comprise possible worlds in the Lewisian sense (Lewis 1986), where each possibility specifies almost everything you can think of.

We usually do not know which of the possibilities in W corresponds to the actual world. Otherwise these possibilities would not be genuine possibilities; our degree of belief function would collapse into the truth value assignment corresponding to the actual world; and you should stop wasting your time by reading this book. All we know is that there is exactly one possibility which is the actual possibility. However, to say that we do not know which possibility is the actual one does not mean that all the possibilities are on a par. Some of them are really far-fetched, while others seem to be more reasonable candidates for the actual possibility.

This gives rise to the following consideration. We can *partition* the set of possibilities, that is, form sets of possibilities that are mutually exclusive (no possibility is in more than one set) and jointly exhaustive (each possibility is in at least one set); and then *order* these sets according to their plausibility. The first set in this ordering contains the possibilities that you take to be the most reasonable candidates for the actual possibility. The second set contains the possibilities which you take to be the second most reasonable candidates. And so on.

If you are still equipped with your possibility distribution from the preceding section you can use your degrees of possibility for the various possibilities to obtain such an *ordered partition*. Note, though, that an ordered partition – in contrast to your possibility distribution – contains no more than ordinal information. While your possibility distribution enables you to say *how* possible you take a possibility to be, an ordered partition only allows you to say that one possibility ω_1 is more plausible than another ω_2. In fact, an ordered partition does not even enable you to express that the difference between your plausibility for ω_1 and for ω_2 is smaller than the difference between your plausibility for ω_2 and for the far-fetched ω_3.

This takes us directly to ranking theory (Spohn 1988), which goes one step further. Rather than merely ordering the possibilities in W, a *pointwise ranking function* $\kappa : W \rightarrow N \cup \{\infty\}$ additionally assigns natural numbers to the sets of possibilities. These numbers represent the degree of disbelief of the various possibilities in W. The result is a *numbered partition* of W,

$$\kappa^{-1}(0), \kappa^{-1}(1), \ldots, \kappa^{-1}(n), \ldots$$

The first member $\kappa^{-1}(0)$ is the set of possibilities which are not disbelieved (which does not mean that they are believed). The second member $\kappa^{-1}(1)$ is the set of possibilities which are disbelieved to degree 1. And so on. It is important to note that many of these sets $\kappa^{-1}(n)$ may be empty, and so do not appear at all in the corresponding ordered partition. (This is where ranking theory goes beyond merely ordered partitions, which are used in form of so called entrenchment orderings in belief revision theory. See Rott's contribution.)

More precisely, a function $\kappa : W \to N \cup \{\infty\}$ from a set of possibilities W into the set of natural numbers extended by ∞, $N \cup \{\infty\}$, is a (normalized) pointwise ranking function just in case $\kappa(\omega) = 0$ for at least one $\omega \in W$, i.e. just in case $\kappa^{-1}(0) \neq \emptyset$. The latter requirement says that you should not disbelieve every possibility, because one possibility is the actual one. A pointwise ranking function $\kappa : W \to N \cup \{\infty\}$ on W induces a *ranking function* $\varrho : \mathcal{A} \to N \cup \{\infty\}$ on a field of propositions $\mathcal{A}$ over W by defining for each $A \in \mathcal{A}$,

$$\varrho(A) = \min\{\kappa(\omega) : \omega \in A\} \quad (= \infty \text{ if } A = \emptyset).$$

This entails that ranking functions $\varrho : \mathcal{A} \to N \cup \{\infty\}$ are (*finitely*) *minimitive*, i.e. for all $A, B \in \mathcal{A}$,

9. $\varrho(A \cup B) = \min\{\varrho(A), \varrho(B)\}$.

As in the case of possibility theory, ranking functions can be directly defined on a field of propositions $\mathcal{A}$ over a set of possibilities W as functions $\varrho : \mathcal{A} \to N \cup \{\infty\}$ such that for all $A, B \in \mathcal{A}$:

$$\varrho(\emptyset) = \infty, \quad \varrho(W) = 0, \quad \varrho(A \cup B) = \min\{\varrho(A), \varrho(B)\}.$$

If ϱ additionally satisfies $\varrho(\mathcal{B}) = \min\{\varrho(A) : A \in \mathcal{B}\}$ for every countable / possibly uncountable $\mathcal{B} \subseteq \mathcal{A}$, ϱ is *countably / completely minimitive*. (Huber 2006a discusses under which conditions ranking functions on fields of propositions are induced by pointwise ranking functions on the set of possibilities one level below.) The conditional ranking function $\varrho(\cdot \mid \cdot) : \mathcal{A} \times \mathcal{A} \to N \cup \{\infty\}$ (based on the unconditional ranking function $\varrho : \mathcal{A} \to N \cup \{\infty\}$) is defined for all $A, B \in \mathcal{A}$ with $A \neq \emptyset$ as

$$\varrho(A \mid B) = \varrho(A \cap B) - \varrho(B),$$

where $\infty - \infty = 0$. Further stipulating $\varrho(\emptyset \mid B) = \infty$ for all $B \in \mathcal{A}$ guarantees that $\varrho(\cdot \mid B) : \mathcal{A} \to N \cup \{\infty\}$ is a ranking function, for every $B \in \mathcal{A}$.

The number $\varrho(A)$ represents the agent's degree of disbelief for the proposition A. If $\varrho(A) > 0$, the agent disbelieves A to a positive degree. Therefore, on pain of inconsistency, she cannot also disbelieve $\overline{A}$ to a positive degree. In other words, at least one of A and $\overline{A}$ has to be assigned rank 0. If $\varrho(A) = 0$, the agent does not disbelieve A to any positive degree. This does not mean, however, that she believes A to a positive degree – the agent may suspend judgement and assign rank 0 to both A and $\overline{A}$. Rather, belief in a proposition is characterized by disbelief in its negation. Thus it is natural to define for each ranking $\varrho : \mathcal{A} \to N \cup \{\infty\}$ a corresponding *belief function* $\beta_\varrho : \mathcal{A} \to Z \cup \{\infty\} \cup \{-\infty\}$ that assigns positive numbers to those propositions that are believed, negative numbers to those that are disbelieved, and 0 to those with respect to which the agent suspends judgement. That is,

$$\beta_\varrho(A) = \varrho\left(\overline{A}\right) - \varrho(A).$$

Each ranking function $\varrho : \mathcal{A} \to N \cup \{\infty\}$ – or equivalently, each belief function $\beta_\varrho : \mathcal{A} \to Z \cup \{\infty\} \cup \{-\infty\}$ – defines a belief set

$$\begin{aligned} B_\varrho &= \left\{A \in \mathcal{A} : \varrho\left(\overline{A}\right) > 0\right\} \\ &= \left\{A \in \mathcal{A} : \beta_\varrho(A) > 0\right\}. \end{aligned}$$

It is the set of all propositions that are believed to some positive degree, or equivalently, whose complements are disbelieved to a positive degree. ϱ's belief set B_ϱ is consistent and deductively closed in the classical finite sense. (The same is true for the set belief set induced by a possibility measure $\Pi : \mathcal{A} \to \Re$, $B_\Pi = \left\{A \in \mathcal{A} : \Pi\left(\overline{A}\right) < 1\right\}$.) Thus ranking theory offers a link between belief and degrees of belief.

As for subjective probabilities there are rules for updating one's epistemic state represented by a ranking function. In case the new information comes in form of a proposition that one strictly speaking learns, ranking theory offers

Update Rule 3 (Plain Conditionalization) *If $\varrho : \mathcal{A} \to N \cup \{\infty\}$ is your ranking function at time t, and between t and t' your learn $E \in \mathcal{A}$ and no logically stronger proposition, then your ranking function at time t' should be $\varrho(\cdot \mid E) : \mathcal{A} \to N \cup \{\infty\}$.*

In case the new information changes your ranks for various propositions, ranking theory offers

Update Rule 4 (Spohn Conditionalization) *If* $\varrho : \mathcal{A} \rightarrow N \cup \{\infty\}$ *is your ranking function at time* t*, and between* t *and* t' *your ranks in the mutually exclusive and jointly exhaustive propositions* $E_1, \ldots, E_n, \ldots$ *(*$E_i \in \mathcal{A}$*) change to* $n_i \in N \cup \{\infty\}$ *with* $\min_i n_i = 0$*, and the finite part of your ranking function does not change on any superset thereof, then your ranking function at time* t' *should be* $\varrho' : \mathcal{A} \rightarrow N \cup \{\infty\}$*, where*

$$\varrho'(\cdot) = \min{}_i \{\varrho(\cdot \mid E_i) + n_i\}.$$

So we see that whenever we substitute 0 for 1, ∞ for 0, $\min$ for $\sum$, $\sum$ for $\prod$, and $>$ for $<$, a true statement about probabilities almost always turns into a true statement about ranking functions. The reason is that each non-standard probability induces a ranking function by letting the ranks be the orders-of-magnitude of the non-standard probability. For more on ranking theory see Spohn's contribution.

One reason why an epistemic agent's degrees of belief should obey the probability calculus is that otherwise she is vulnerabe to a Dutch Book. For similar reasons she should update her subjective probability according to strict or Jeffrey Conditionalization, depending on the format of the new information. Why should degrees of disbelief obey the ranking calculus? And why should an epistemic agent update her ranking function according to plain or Spohn Conditionalization? The answers to these questions require a bit of terminology. A set of propositions $\mathcal{B} \subseteq \mathcal{A}$ is *consistent in the finite / countable / complete sense* if and only if for every finite / countable / possibly uncountable $B \subseteq \mathcal{B}$: $\bigcap B \neq \emptyset$. $\mathcal{B}$ is *deductively closed in the finite / countable / complete* sense if and only if for every finite / countable / possibly uncountable $B \subseteq \mathcal{B}$ and every $A \in \mathcal{A}$: if $\bigcap B \subseteq A$, then $A \in \mathcal{B}$.

Furthermore, an epistemic agent's *degree of entrenchment* for a proposition A is the number of independent and minimally positively reliable information sources it takes for the agent to give up her belief that $\overline{A}$. If the agent does not believe $\overline{A}$ to begin with, her degree of entrenchment for A is 0. If no finite number of information sources is able to make the agent give up her belief that $\overline{A}$, her degree of entrenchement for A is ∞. Let us suppose Sophia believes that Sacramento is not the capital of California. Her degree of entrenchment for the proposition that Sacramento is the capital of California can be measured by putting her on Sunset Boulevard and count the number of people passing by and telling her that Sacramento is the capital of California. Degrees of entrenchment are similarly related to ranks as are fair betting ratios to subjective probabilities.

With this in mind the first thing we can say is that the belief set B_Π of a possibility measure $\Pi : \mathcal{A} \rightarrow \Re$ is not consistent and deductively closed in the countable (and hence complete) sense – witness the example on page 13. The second thing this enables us to do is to answer the question why degrees of disbelief should obey the ranking calculus. They should do so, because an agent's belief set is and will always be consistent and deductively closed in the finite / countable / complete sense just in case her entrenchment function is a finitely / countably / completely minimitive ranking function and, depending on the format of the evidence, she updates according to plain or Spohn conditionalization. The third point to add is that this is the very same reason why ranking functions should be updated by plain or Spohn Conditionalization: Every possible current or future belief set $B_{\varrho'}$ based on the entrenchment function ϱ' is consistent and deductively closed in the finite / countable / complete sense just in case ϱ' results from the finitely / countably / completely minimitive ranking function ϱ by plain or Spohn Conditionalization, depending on whether the agent strictly speaking learns a proposition or merely changes her ranks on a partition. See Huber (2007).

5 Belief Revision and Nonmonotonic Reasoning

5.1 Belief and Belief Revision

In the meantime we have moved from degrees of belief to belief, and found ranking theory to provide a link between these two notions. While some people hold the view that degrees of belief are more basic than beliefs simpliciter, others adopt the opposite view. This is generally true of traditional epistemology which is mainly concerned with the notion of knowledge and its tripartite definition as justified true belief. Belief in this sense comes in three "degrees": The ideal or rational epistemic agent either believes A, or else she believes $\overline{A}$, or else she believes neither A nor $\overline{A}$, but suspends judgement. Ordinary epistemic agents sometimes believe both A and $\overline{A}$, but we will assume that they should not do so, and hence ignore this case.

An agent's epistemic state is thus characterized by the set of propositions she believes, her *belief set*. Such a belief set is required to be consistent and deductively closed (Hintikka 1961). Here a belief set is usually represented as a set of sentences from a language $\mathcal{L}$ rather than as a set of propositions. The question adressed by *belief revision theory* (Alchourrón & Gärdenfors & Makinson 1985, Gärdenfors 1988, Gärdenfors & Rott 1995) is how an ideal or rational epistemic agent should revise her belief set B if she learns new information in form of a proposition α. If α is consistent with B in the syntactical sense that B does not logically imply the negation of α, $B \nvdash \neg\alpha$, the agent should simply add α to B and close this set under (classical) logical consequence. In more technical terms, her new belief set is the set of logical consequences of $B \cup \{\alpha\}$,

$$Cn\left(B \cup \{\alpha\}\right) = \{\beta \in \mathcal{L} : B \cup \{\alpha\} \vdash \beta\}.$$

The interesting case is when the new information α contradicts the old belief set B in the sense that $B \vdash \neg\alpha$. Here the basic idea is that the agent's new belief set should contain the new information α and as many of the old beliefs in B as is allowed by requiring the new belief set to be consistent and deductively closed. To state this more precisely, let us introduce the notion of a *contraction*. To contract a statement α from a belief set B is to give up the belief that α is true, but to keep as many of the remaining beliefs from B while ensuring consistency and deductive closure. With $B \dot{-} \alpha$ as the agent's new belief set after contracting her old belief set B by α, the A(lchourrón)G(ärdenfors)M(akinson) postulates for contraction $\dot{-}$ can be stated as follows. For every set of sentences $B \subseteq \mathcal{L}$ and any sentences $\alpha, \beta \in \mathcal{L}$:

$\dot{-}$1. If $B = Cn\left(B\right)$, then $B \dot{-} \alpha = Cn\left(B \dot{-} \alpha\right)$. Deductive Closure

$\dot{-}$2. $B \dot{-} \alpha \subseteq B$. Inclusion

$\dot{-}$3. If $\alpha \notin Cn\left(B\right)$, then $B \dot{-} \alpha = B$. Vacuity

$\dot{-}$4. If $\alpha \notin Cn\left(\emptyset\right)$, then $\alpha \notin Cn\left(B \dot{-} \alpha\right)$. Success

$\dot{-}$5. If $Cn\left(\{\alpha\}\right) = Cn\left(\{\beta\}\right)$, then $B \dot{-} \alpha = B \dot{-} \beta$. Preservation

$\dot{-}$6. If $B = Cn\left(B\right)$, then $B \subseteq Cn\left(\left(B \dot{-} \alpha\right) \cup \{\alpha\}\right)$. Recovery

$\dot{-}$7. If $B = Cn\left(B\right)$, then $\left(B \dot{-} \alpha\right) \cap \left(B \dot{-} \beta\right) \subseteq B \dot{-} \left(\alpha \wedge \beta\right)$.

$\dot{-}$8. If $B = Cn\left(B\right)$ and $\alpha \notin B \dot{-} \left(\alpha \wedge \beta\right)$, then $B \dot{-} \left(\alpha \wedge \beta\right) \subseteq B \dot{-} \alpha$.

$\dot{-}$1 says that the contraction of B by α, $B \dot{-} \alpha$, should be deductively closed, if B is deductively closed. $\dot{-}$2 says that a contraction should not give rise to new beliefs not previously held. $\dot{-}$3 says that the epistemic agent should not change her beliefs when she gives up a sentence she does not believe to begin with. $\dot{-}$4 says that, unless α is tautological, the agent should indeed give up her belief that α is true if she contracts by α. $\dot{-}$5 says that the particular formulation of the sentence you give up should not matter; in other words, the objects of belief really shoud be propositions rather than sentences. $\dot{-}6$ says the agent should recover her old beliefs if she adds a sentence she just gave up. According to $\dot{-}7$ the agent should not give up more beliefs when contracting by $\alpha \wedge \beta$ than the ones she gives up when she contracts by α or when she contracts by β. $\dot{-}8$ finally requires the agent not to give up more beliefs than necessary: If the agent gives up α when she contracts by $\alpha \wedge \beta$, she should not give up more than she gives up when contracting by α alone. Rott (2001) discusses many further principles.

Given the notion of a contraction we can now state what the agent's new belief set should look like when B is her old belief set and she gets the new information α. First, the agent should clear B to make it consistent with α. That is, the agent first should contract B by $\neg\alpha$. Then she should simply add α to this set and close under (classical) logical consequence. This gives us the agent's new belief set $B \dot{+} \alpha$, her old belief set B revised by α. The recipe just described is known as the *Levi identity*:

$$B \dot{+} \alpha = Cn\left(\left(B \dot{-} \neg\alpha\right) \cup \{\alpha\}\right).$$

Revision $\dot{+}$ defined in this way satisfies a corresponding list of properties. For every set of sentences $B \subseteq \mathcal{L}$ and any sentences $\alpha, \beta \in \mathcal{L}$:

$\dot{+}1$. $B \dot{+} \alpha = Cn\left(B \dot{+} \alpha\right)$.

$\dot{+}2$. $\alpha \in B \dot{+} \alpha$.

$\dot{+}3$. If $\neg\alpha \notin Cn(B)$, then $B \dot{+} \alpha = Cn(B \cup \{\alpha\})$.

$\dot{+}4$. If $\neg\alpha \notin Cn(\emptyset)$, then $\perp \notin B \dot{+} \alpha$.

$\dot{+}5$. If $Cn(\{\alpha\}) = Cn(\{\beta\})$, then $B \dot{+} \alpha = B \dot{+} \beta$.

$\dot{+}6$. If $B = Cn(B)$, then $\left(B \dot{+} \alpha\right) \cap B = B \dot{-} \neg\alpha$.

$\dot{+}7$. If $B = Cn(B)$, then $B \dot{+} (\alpha \wedge \beta) \subseteq Cn\left(\left(B \dot{+} \alpha\right) \cup \{\beta\}\right)$.

$\dot{+}8$. If $B = Cn(B)$ and $\neg\beta \notin B \dot{+} \alpha$, then $Cn\left(\left(B \dot{+} \alpha\right) \cup \{\beta\}\right) \subseteq B \dot{+} (\alpha \wedge \beta)$.

In standard belief revision theory the new information is always part of the new belief set. *Non-prioritized belief revision* relaxes this requirement (Hansson 1999). The epistemic agent might consider the new information to be too implausible to be added and decide to reject it; or she might add only a sufficiently plausible part of the new information; or else, she might add the new information and then check for consistency, which makes her give up part or all of the new information again, because her old beliefs turn out to be more *entrenched*.

The latter notion provides the connection to degrees of belief. In order to decide which part of her belief set she wants to give up, belief revision theory equips the epistemic agent with an *entrenchment ordering*. Technically, these are relations $\preceq$ on $\mathcal{L}$ such that for all $\alpha, \beta, \gamma \in \mathcal{L}$:

E1. $\alpha \preceq \beta, \quad \beta \preceq \gamma \quad \Rightarrow \quad \alpha \preceq \gamma$ — Transitivity

E2. $\alpha \vdash \beta \quad \Rightarrow \quad \alpha \preceq \beta$ — Dominance

E3. $\alpha \preceq \alpha \wedge \beta$ or $\beta \preceq \alpha \wedge \beta$ — Conjunctivity

E4. $\bot \notin Cn\,(B) \quad \Rightarrow \quad [\alpha \notin B \quad \Leftrightarrow \quad \forall \beta \in \mathcal{L} : \alpha \preceq \beta]$ — Minimality

E5. $\forall \alpha \in \mathcal{L} : \alpha \preceq \beta \quad \Rightarrow \quad \beta \in Cn\,(\emptyset)$ — Maximality

B is a fixed set of background beliefs. Given an entrenchment ordering $\preceq$ on $\mathcal{L}$ we can define a revision $\dot{+}$ as follows:

$$B \dot{+} \alpha = \{\beta \in B : \neg\alpha \prec \beta\} \cup \{\alpha\}\,,$$

where $\alpha \prec \beta$ holds just in case $\alpha \preceq \beta$ and $\beta \npreceq \alpha$. Then one can prove the following *representation theorem*:

Theorem 1 *Each entrenchment ordering $\preceq$ on $\mathcal{L}$ induces a revision operator $\dot{+}$ on $\mathcal{L}$ satisfying $\dot{+}$1-8. For each revision operator $\dot{+}$ on $\mathcal{L}$ satisfying $\dot{+}$1-8 there is an entrenchment ordering $\preceq$ on $\mathcal{L}$ that induces $\dot{+}$.*

It is, however, fair to say that belief revision theorists distinguish between degrees of belief and entrenchment. Entrenchment, so they say, characterizes the agent's unwillingness to give up a particular belief, which may be different from her degree of belief for the respective sentence or proposition. Although this distinction seems to violate Occam's razor by unnecessarily introducing an additional epistemic level, it corresponds to Spohn's parallelism between subjective probabilities and ranking functions. The reader will also be reminded of the distinction between the credal and the pignistic level in the transferable belief model, as well as the belief/degree of belief view in connection with subjective probability theory.

5.2 Belief and Nonmonotonic Reasoning

A premise β classically entails a conclusion γ, $\beta \vdash \gamma$, just in case γ is true in every model or truth value assignment in which β is true. The classical consequence relation $\vdash$ (conceived of as a relation between two sentences rather than as a relation between a set of sentences, the premises, and a sentence, the conclusion) is non-ampliative in the sense that the conclusion of a classically valid argument does not convey information that goes beyond what is contained in the premise.

It has the following *monotonicity* property. For any sentences $\alpha, \beta, \gamma \in \mathcal{L}$,

$$\beta \vdash \gamma, \quad \alpha \vdash \beta \quad \Rightarrow \quad \alpha \vdash \gamma.$$

That is, if γ follows from β, then γ follows from any logically stronger sentence α. However, everyday reasoning is ampliative. When Sophia sees the thermometer at 33° Celsius she infers that it is not too cold to wear her sundress. If Sophia additionally sees that the thermometer is placed above the oven where she is boiling her pasta, she will not infer so anymore. Although classically invalid, these are reasonable inferences. *Nonmonotonic reasoning* is the study of such reasonable consequence relations which violate monotonicity.

For a fixed set of background beliefs B, the revision operators $\dot{+}$ from the previous section give rise to nonmonotonic consequence relations $|\!\sim$ as follows (Makinson & Gärdenfors 1991):

$$\alpha \mid\!\sim \beta \quad \Leftrightarrow \quad \beta \in B \dot{+} \alpha.$$

Nonmonotonic consequence relations on a language $\mathcal{L}$ are supposed to satisfy the following principles from Kraus & Lehmann & Magidor (1990) (for an overview see Makinson 1994).

KLM1. $\alpha \mid\!\sim \alpha$ — Reflexivity

KLM2. $\vdash \alpha \leftrightarrow \beta, \quad \alpha \mid\!\sim \gamma \quad \Rightarrow \quad \beta \mid\!\sim \gamma$ — Left Logical Equivalence

KLM3. $\vdash \alpha \rightarrow \beta, \quad \gamma \mid\!\sim \alpha \quad \Rightarrow \quad \gamma \mid\!\sim \beta$ — Right Weakening

KLM4. $\alpha \wedge \beta \mid\!\sim \gamma, \quad \alpha \mid\!\sim \beta \quad \Rightarrow \quad \alpha \mid\!\sim \gamma$ — Cut

KLM5. $\alpha \mid\!\sim \beta, \quad \alpha \mid\!\sim \gamma \quad \Rightarrow \quad \alpha \wedge \beta \mid\!\sim \gamma$ — Cautious Monotonicity

KLM6. $\alpha \mid\!\sim \beta, \quad \alpha \mid\!\sim \gamma \quad \Rightarrow \quad \alpha \vee \beta \mid\!\sim \gamma$ — Or

The standard interpretation of a nonmonotonic consequence relation $|\!\sim$ is "If ..., normally ...". Normality among worlds is spelt out in terms of *preferential models* $\langle S, l, \prec \rangle$ for $\mathcal{L}$, where S is a set of states, and $l : S \to Mod_{\mathcal{L}}$ is a function from S to the set of models for $\mathcal{L}$ that assigns each state s its world $l\,(s)$. The abnormality relation $\prec$ is a strict partial order on $Mod_{\mathcal{L}}$ that satisfies a certain smoothness condition. For our purposes it suffices to note that the order among the worlds that is induced by a pointwise ranking functions is such an abnormality relation. Given a preferential model $\langle S, l, \prec \rangle$ we can define a nonmonotonic consequence relation $|\!\sim$ as follows:

$$\alpha \mid\!\sim \beta \quad \Rightarrow \quad \forall s \in \widehat{\alpha}\,(\forall t \in \widehat{\alpha} : t \not\prec s \to l\,(s) \models \beta)\,,$$

where $\widehat{\alpha} = \{s \in S : l\,(s) \models \alpha\}$ is the set of states in whose worlds α is true. That is, $\alpha \mid\!\sim \beta$ holds just in case β is true in the least abnormal among the α worlds. Then one can show the following representation theorem.

Theorem 2 *Each preferential model $\langle S, l, \prec \rangle$ for $\mathcal{L}$ induces a nonmonotonic consequence relation $|\!\sim$ on $\mathcal{L}$ satisfying KLM1-6. For each nonmonotonic consequence relation on $\mathcal{L}$ satisfying KLM1-6 there is a preferential model $\langle S, l, \prec \rangle$ for $\mathcal{L}$ that induces $|\!\sim$.*

So whereas the classical consequence relation preserves truth in all logically possible worlds, nonmonotonic consequence relations preserve truth in all least abnormal worlds. For a different semantics in terms of inhibition nets see Leitgeb (2004).

What is of particular interest to us is the fact that these nonmonotonic consequence relations can be induced by a fixed set of background beliefs B and various forms of degrees of belief over B. We will not attempt to indicate how this works. Makinson's contribution is an excellent presentation of ideas underlying nonmonotonic reasoning and its relation to degrees of belief.

Acknowledgements

I am very grateful to Christoph Schmidt-Petri for many helpful comments.

My research was supported by the Ahmanson Foundation as well as the Alexander von Humboldt Foundation, the Federal Ministry of Education and Research, and the Program for the Investment in the Future (ZIP) of the German Government through a Sofja Kovalevskaja Award to Luc Bovens, while I was member of the *Philosophy, Probability, and Modeling* group at the Center for Junior Research Fellows at the University of Konstanz.

References

[1] Alchourrón, Carlos E. & Gärdenfors, Peter & Makinson, David (1985), *On the Logic of Theory Change: Partial Meet Contraction and Revision Functions*. *Journal of Symbolic Logic* **50**, 510-530.

[2] Carnap, Rudolf (1962), *Logical Foundations of Probability*. 2nd ed. Chicago: University of Chicago Press.

[3] Cox, Richard T. (1961), *The Algebra of Probable Inference*. Baltimore, MD: Johns Hopkins Press.

[4] Dempster, Arthur P. (1968), A Generalization of Bayesian Inference. *Journal of the Royal Statistical Society*. *Series B (Methodological)* **30**, 205-247.

[5] Dubois, Didier & Prade, Henri (1980), *Possibility Theory. An Approach to Computerized Processing of Uncertainty*. New York: Plenum.

[6] Foley, Richard (1992), The Epistemology of Belief and the Epistemology od Degrees of Belief. *American Philosophical Quaterly* **29**, 111-121.

[7] Frankish, Keith (2004), *Mind and Supermind*. Cambridge: Cambridge University Press.

[8] Gärdenfors, Peter (1988), *Knowledge in Flux. Modeling the Dynamics of Epistemic States*. Cambridge, MA: MIT Press.

[9] Gärdenfors, Peter & Rott, Hans (1995), Belief Revision. In D.M. Gabbay & C.J. Hogger & J.A. Robinson (eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming*. Vol. 4. *Epistemic and Temporal Reasoning*. Oxford: Clarendon Press, 35-132.

[10] Haenni, Rolf & Lehmann, Norbert (2003), Probabilistic Argumentation Systems: A New Perspective on Dempster-Shafer Theory. *International Journal of Intelligent Systems* **18**, 93-106.

[11] Hájek, Alan (2003), What Conditional Probability Could Not Be. *Synthese* **137**, 273-323.

[12] —— (2005), Scotching Dutch Books? Philosopical Perspectives 19 (Epistemology), 139-151.

[13] Halpern, Joseph Y. (2003), *Reasoning About Uncertainty*. Cambridge, MA: MIT Press.

[14] Hansson, Sven Ove (1999), A Survey of Non-Prioritized Belief Revision. *Erkenntnis* **50**, 413-427.

[15] Hawthorne, James & Bovens, Luc (1999), The *Preface*, the *Lottery*, and the Logic of Belief. *Mind* **108**, 241-264.

[16] Hempel, Carl Gustav (1962), Deductive-Nomological vs. Statistical Explanation. In H. Feigl & G. Maxwell (eds.), *Scientific Explanation, Space and Time*. *Minnesota Studies in the Philosophy of Science* 3. Minneapolis: University of Minnesota Press, 98-169.

[17] Hintikka, Jaakko (1961), *Knowledge and Belief. An Introduction to the Logic of the Two Notions*. Ithaca, NY: Cornell University Press.

[18] Huber, Franz (2006a), Ranking Functions and Rankings on Languages. *Artificial Intelligence* **170**, 462-471 .

[19] —— (2006b), Confirmation. To appear in J. Fieser & B. Dowdon (eds.), *The Internet Encyclopedia of Philosophy*.

[20] —— (2007), The Consistency Argument for Ranking Functions. To appear in *Studia Logica*.

[21] Jeffrey, Richard C. (1970), Dracula Meets Wolfman: Acceptance vs. Partial Belief. In M. Swain (ed.), *Induction, Acceptance, and Rational Belief*. Dordrecht: Reidel, 157-185.

[22] —— (2004), *Subjective Probability: The Real Thing*. Cambridge: Cambridge University Press.

[23] Joyce, James M. (1998), A Nonpragmatic Vindication of Probabilism. *Philosophy of Science* **65**, 575-603.

[24] Kolmogorov, Andrej N. (1956), *Foundations of the Theory of Probability*. 2nd ed. New York: Chelsea Publishing Company.

[25] Krantz, David H. & Luce, Duncan R. & Suppes, Patrick & Tversky, Amos (1971), *Foundations of Measurement*. Vol. I. New York: Academic Press.

[26] Kraus, Sarit & Lehmann, Daniel & Magidor, Menachem (1990), Nonmonotonic Reasoning, Preferential Models, and Cumulative Logics. *Artificial Intelligence* **40**, 167-207.

[27] Kyburg, Henry E. Jr. (1961), *Probability and the Logic of Rational Belief*. Middletown, CT: Wesleyan University Press.

[28] —— (1970), Conjunctivitis. In M. Swain (ed.), *Induction, Acceptance, and Rational Belief*. Dordrecht: Reidel, 232-254.

[29] Leitgeb, Hannes (2004), *Inference on the Low Level*. Dordrecht: Kluwer.

[30] Levi, Isaac (1980), *The Enterprise of Knowledge*. Cambrdige, MA: MIT Press.

[31] Lewis, David (1980), A Subjectivist's Guide to Objective Chance. In R.C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*. Vol. II. Berkeley: University of Berkeley Press, 263-293. Reprinted in D. Lewis (1986), *Philosophical Papers*. Vol. II. Oxford: Oxford University Press, 83-113.

[32] Lewis, David (1986), *On the Plurality of Worlds*. Oxford: Blackwell.

[33] Makinson, David (1965), The Paradox of the Preface. *Analysis* **25**, 205-207.

[34] —— (1994), General Patterns in Nonmonotonic Reasoning. In D.M. Gabbay & C.J. Hogger & J.A. Robinson (eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming*. Vol. 3. *Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford: Clarendon Press, 35-110.

[35] Makinson, David & Gärdenfors, Peter (1991), Relations between the Logic of Theory Change and Nonmonotonic Logic. A. Fuhrmann & M. Morreau (eds.), *The Logic of Theory Change*. Berlin: Springer, 185-205.

[36] Roorda, Jonathan (1995), Revenge of Wolfman: A Probabilistic Explication of Full Belief. http://www.princeton.edu/˜bayesway/pu/Wolfman.pdf

[37] Rott, Hans (2001), *Change, Choice, and Inference. A Study of Belief Revision and Nonmonotonic Reasoning*. Oxford: Oxford University Press.

[38] Savage, Leonard J. (1972). *The Foundations of Statistics*. 2nd ed. New York: Dover.

[39] Shafer, Glenn (1976), *A Mathematical Theory of Evidence*. Princteton, NJ: Princeton University Press.

[40] Skyrms, Brian (1987), Dynamic Coherence and Probability Kinematics. *Philosophy of Science* **54**, 1-20.

[41] Smets, Philippe (2002), Showing Why Measures of Quantified Beliefs are Belief Functions. In B. Bouchon & L. Foulloy & R.R. Yager (eds.), Intelligent Systems for Information Processing: From Representation to Applications. Amsterdam: Elsevier, 265-276.

[42] Smets, Philippe & Kennes, Robert (1994), The Transferable Belief Model. *Artifical Intelligence* **66**, 191-234.

[43] Spohn, Wolfgang (1988), Ordinal Conditional Functions: A Dynamic Theory of Epistemic States. In W.L. Harper & B. Skyrms (eds.), *Causation in Decision, Belief Change, and Statistics II*. Dordrecht: Kluwer, 105-134.

[44] —— (1997), The Intentional versus the Propositional Conception of the Objects of Belief. In L. Villegas & M. Rivas Monroy & C. Martinez (eds.), *Proceedings of the Congress on Truth, Logic, and Representation of the World*. Santiago de Compostela, 266-286.

[45] Stalnaker, Robert C. (1984), *Inquiry*. Cambridge, MA: MIT Press.

[46] Walley, Peter (1991), *Statistical Reasoning With Imprecise Probabilities*. New York: Chapman and Hall.

[47] Zadeh, Lotfi A. (1978), Fuzzy Sets as a Basis for a Theory of Possibility. *Fuzzy Sets and Systems* **1**, 3-28.