28

Degree of Belief is Expected Truth Value

Nicholas J. J. Smith

This chapter presents a solution to a problem engendered by the following two claims:

(A) Vagueness gives rise to degrees of belief.

(B) These degrees of belief do not behave in the same ways as degrees of belief arising from uncertainty: they do not conform to the laws of probability.

The problem is to give a clear account of the relationship between degrees of belief and subjective probabilities. The solution to be presented here also involves degrees of truth: in outline, the proposal is that one's degree of belief in a proposition *P* is one's expectation of *P*'s degree of truth. Those who already believe that vagueness should be handled using degrees of truth will believe (A) and (B). So the chapter can be read as solving a problem which arises for degree theorists. It can also be read as providing a positive argument in favour of degrees of truth, directed at those who do not start out believing that vagueness should be handled using degrees of truth, but do start out believing (A) and (B): the argument is that the best solution to the problem engendered by (A) and (B) employs degrees of truth.

28.1 VAGUENESS-BASED AND UNCERTAINTY-BASED DEGREES OF BELIEF

Suppose we have a Sorites series leading from tall men down to short men. Suppose also that we have accepted a degree-theoretic account of vagueness—so we think that

An earlier version of this chapter—including a surrounding discussion of how degrees of truth can be incorporated into the framework of Stalnakerian pragmatics—was presented at the Arché Vagueness Conference in St Andrews on 8 June 2007 (see Smith (2008, §5.3) for that surrounding discussion). Other earlier versions were presented at the Annual Conference of the Australasian Association of Philosophy on 5 July 2007, in the Higher Seminar in Theoretical Philosophy in the Department of Philosophy at Lund University on 1 April 2008, and in the Current Projects Seminar in the Centre for Time at the University of Sydney on 25 August 2008. Thanks to the audiences on those occasions for useful feedback. For helpful discussions, I am grateful to Staffan Angere, Richard Dietz, Andy Egan, Michael McDermott, Peter Milne, Josh Parsons, Wlodek Rabinowicz, and Roy Sorensen. Thanks also to the two anonymous referees for their comments and to the Australian Research Council for research support.

'This man is tall' goes gradually from 1 true, said of men at the beginning of the series, down to 0 true, said of men at the end. Then what attitude should we adopt to (the proposition expressed by) 'This man is tall' as we consider various men in the series? Surely we should go from being fully committed to the proposition at the beginning of the series, to fully rejecting it by the end of the series, via a gradually changing series of intermediate states of partial belief, which decrease in degree of confidence as we progress down the series.[1] So degree theorists should certainly accept (A). But it seems that non-degree theorists should accept (A) too. Consider Schiffer (2000, 223–4):

> Sally is a rational speaker of English, and we're going to monitor her belief states throughout the following experiment. Tom Cruise, a paradigmatically non-bald person, has consented, for the sake of philosophy, to have his hairs plucked from his scalp one by one until none are left. Sally is to witness this, and will judge Tom's baldness after each plucking. The conditions for making baldness judgments—lighting conditions, exposure to the hair situation on Tom's scalp, Sally's sobriety and perceptual faculties, etc.—are ideal and known by Sally to be such . . . Let the plucking begin.
>
> Sally starts out judging with absolute certainty that Tom is not bald; that is, she believes to degree 1 that Tom is not bald and to degree 0 that he is bald. This state of affairs persists through quite a few pluckings. At some point, however, Sally's judgment that Tom isn't bald will have an ever-so-slightly-diminished confidence, reflecting that she believes Tom not to be bald to some degree barely less than 1. The plucking continues and as it does the degree to which she believes Tom not to be bald diminishes while the degree to which she believes him to be bald increases . . . Sally's degrees of belief that Tom is bald will gradually increase as the plucking continues, until she believes to degree 1 that he is bald.
>
> Although I'll have a little more to say about this later, for now I'm going to assume that the qualified judgments about Tom's baldness that Sally would make throughout the plucking express partial beliefs. After all, the hallmark of partial belief is qualified assertion, and, once she was removed from her ability to make unqualified assertions, Sally would make qualified assertions in response to queries about Tom's baldness.

Other things that we might say about the case—things that would avoid claiming that Sally has degrees of belief—are (i) that Sally fully believes that Tom is not bald until a particular hair is removed, from which point on she fully believes he is bald; (ii) that Sally fully believes that Tom is not bald until a particular hair is removed, at which point she enters an indeterminate state in which she does not believe (to any degree, even 0) that Tom is not bald and does not believe (to any degree, even 0) that Tom is bald, and then when another particular hair is removed Sally comes to fully believe that Tom is bald; and (iii) that Sally does not have attitudes towards propositions such as 'Tom is bald', but only towards propositions such as 'Tom is bald to degree *x*' or ' "Tom is bald" is true to degree *x*', each of which she either fully believes or fully rejects. The problem with these approaches is that they do not fit the phenomena. Contra (i) and (iii), Sally certainly seems to be unsure as to what to believe and say about Tom's baldness, at various points in the process, and contra (ii), she does not have one catch-all 'confused state', which she enters, remains in, then leaves: rather, she seems clearly to become *less and less sure* that Tom is not bald, and then later *more and more sure* that he is.

[1] I treat the terms 'degree of belief', 'partial belief' and 'credence' as synonyms.

The proponent of (iii) may reply that Sally's qualified assertion that Tom is bald—behaviour which *seems* clearly to indicate that there is some *P* such that Sally is unsure whether *P*—should in fact be understood as a full-on assertion that Tom is bald to an intermediate degree. But this response is rather implausible on the face of it, and furthermore (iii) involves a strange separation between truth on the one hand, and belief and assertion on the other. The view involves a semantics which assigns degrees of truth to atomic propositions such as 'Tom is bald', but tells us that we cannot believe or assert such propositions. Rather, we must believe and assert meta-level propositions of the form ' "Tom is bald" is true to degree *x*', or propositions about degrees, such as 'Tom's degree of baldness is *x*'. This kind of separation (between truth on the one hand, and belief and assertion on the other) should be regarded as a last resort, to be considered only if it were shown that we cannot, for some reason, adopt what should be the default position, namely that the very same things both have truth values and are the contents of beliefs and assertions.[2]

So we need to countenance degrees of belief arising from vagueness. Doing so will not cause us any problem, however, if these degrees of belief are just the same as the kind with which we are already familiar: the kind that arise from *uncertainty* about the truth of propositions (in cases not involving vagueness), and are handled formally by means of probability theory. However it seems that degrees of belief arising from vagueness do *not* behave in the same ways as degrees of belief arising from uncertainty. To adapt and augment an example of Schiffer's: Suppose that Sally is about to meet her long-lost brother Sali. She has been told that he is either very tall or very short, but she has no idea which (so she does know that he is not a borderline case), and she has been told that he is either hirsute or totally bald, but she has no idea which (so she does know that he is not a borderline case). As a result of her uncertainty, she believes both of the propositions 'Sali is tall' and 'Sali is bald' to degree 0.5. Suppose also that Sally regards these two propositions as independent: supposing one to be true would have no bearing on her beliefs about the other. Then, for familiar reasons, she should believe 'Sali is tall and bald' to degree 0.25. Now suppose that mid-way through Schiffer's experiment, when Sally's degree of belief that Tom is bald is 0.5, she also believes to degree 0.5 that Tom is tall—on the basis of looking at him and seeing that he is a classic borderline case of tallness.[3] Then what should be her degree of belief that Tom is tall and bald? The answer 0.5 suggests itself very strongly: certainly the answer 0.25 seems wrong. If you don't think so, then just add more conjuncts (e.g. funny, nice, intelligent, cool, old—where Sally knows of Sali only that he is not a borderline case of any of them, and of Tom that he is a classic borderline case of all of them): the more independent conjuncts you add, the lower the uncertainty-based degree of belief should go, but this does not seem to be the case for the vagueness-based degree of belief (Schiffer, 2000, 225), (MacFarlane, 2006, 6).

So it seems that (B) is true, as well as (A). This means that we must abandon the familiar *identification* of degrees of belief with subjective probabilities, and offer a new

[2] John MacFarlane's view (this volume) suffers from this problem.
[3] Suppose, for the sake of the example, that Tom Cruise is borderline tall.

account of their relationship. In Section 28.2 I critique one kind of account. In Section 28.3 I present my own view, and in Section 28.4 I reply to objections to this view.

28.2 TWO KINDS OF DEGREE OF BELIEF?

One thought in response to (A) and (B) is that there are *two kinds* of degree of belief: uncertainty-based degrees of belief and vagueness-based degrees of belief. Schiffer holds a view of this sort. He distinguishes SPB's ('standard partial beliefs') and VPB's ('vagueness-related partial beliefs'). In his view, we have two distinct systems of degrees of belief: an assignment of SPB's to propositions, which obey the laws of probability, and an assignment of VPB's to propositions, which obey the laws of standard fuzzy propositional logic (i.e. $\text{VPB}(\neg p) = 1 - \text{VPB}(p), \text{VPB}(p \wedge q) = \min\{\text{VPB}(p), \text{VPB}(q)\}$ and $\text{VPB}(p \vee q) = \max\{\text{VPB}(p), \text{VPB}(q)\}$).

There is a grave problem for any proposal which posits two different systems of degrees of belief, where it is allowed that a subject may have a degree of belief of one kind of strength n in a proposition P, and a degree of belief of another kind of strength $m \neq n$ in the same proposition P. The problem is that the very idea of degree of belief is made sense of via the thought that a degree of belief that P is a strength of tendency to act as if P. As Ramsey (1990, 65–6) puts it:

> the degree of a belief is a causal property of it, which we can express vaguely as the extent to which we are prepared to act on it. ... it is not asserted that a belief is an idea which does actually lead to action, but one which would lead to action in suitable circumstances ... The difference [between believing more firmly and believing less firmly] seems to me to lie in how far we should act on these beliefs.

But one simply cannot have two different strengths of tendency to act as if P, in a given set of circumstances. Consider, for example, the proposition that Fido is dangerous. When Fido enters the room, one will do some particular thing, for example sit still, or jump and run. When Fido looks at one, one will do some particular thing, for example tremble, or offer him some beef jerky. When Fido barks, one will do some particular thing, for example scream; and so on. One cannot both back away slowly *and* run screaming (at the same time), and it cannot both take Fido getting within two metres of one to make one run away, *and* require Fido getting within one metre to make one run. So one cannot both tend strongly to act as if Fido is dangerous, *and* tend weakly to act as if Fido is dangerous—at least not if there is to be any sort of transparent relationship between these tendencies and the way one actually acts. But given that a degree of belief just is a strength of tendency to act, this means that one cannot have two different degrees of belief in the same proposition.

The proponent of two kinds of degrees of belief might offer a number of responses here. (1) She might deny that there *is* a transparent relationship between tendencies to act and the way one actually acts. So, in the case of Fido, one might have both a strong tendency to act as if Fido is dangerous, and a weak tendency, and these interact so as to make one behave in particular ways in particular situations (ways that we would like to describe as indicating that one has a mid-strength tendency to act as if Fido is dangerous—although on the current proposal, we cannot straightforwardly say this). But for this view to get off the ground, we would need to be told exactly *how* degrees of belief of the two sorts combine to produce certain behaviour, and furthermore, the view threatens to make it impossible for us ever to know (even roughly) someone's degree(s) of belief in a given proposition. (2) She might say that although there are indeed two kinds of degrees of belief, they always have the same strength, for every proposition. But clearly this would run us headlong into the problem discussed above, that partial beliefs arising from vagueness do not and should not behave in the same ways as partial beliefs arising from uncertainty. (3) She might deny that degrees of belief are to be understood in terms of strength of tendency to act. But any view which disconnects degree of belief from tendency to act threatens to undermine the utility of the notion of degree of belief, and furthermore any candidate replacement proposal—for example, the view that the difference between believing more firmly and believing less firmly is a matter of *strength of feeling*[4]—would seem to face the very same problem (one cannot have two different intensities of feeling about one proposition). (4) She might claim that one never has both kinds of degree of belief in the same proposition at the same time. For suppose, for reductio, that you have an uncertainty-based degree of belief of 0.3 that Dobbin wins the race, and a vagueness-based degree of belief of 0.5 that Dobbin wins the race. How could you have acquired *both* these beliefs? In order to acquire the first, you would need to lack evidence concerning who wins. In order to acquire the second, you would need to have all the relevant evidence, and see that it—i.e. the world itself—leaves it unsettled who wins.[5] So clearly you could not have both these degrees of belief at once. There are still problems for this view, however. First, we need to be told how to reason with several propositions—and compounds thereof—in some of which we have degrees of belief of one type, and in others of which we have degrees of belief of the other type. Second, what justifies saying that we have here two non-interacting systems of degrees of belief, rather than one system, which assigns degrees to all propositions, but where these degrees behave differently in different situations (e.g. sometimes they obey the laws of probability, sometimes they do not)?

This is the remaining possibility regarding the relationship between vagueness-based degrees of belief, and uncertainty-based degrees of belief: the suggestion that what we have is one univocal notion of degree of belief—one single system of assignments of degrees of belief to propositions—but where the degrees assigned sometimes behave in accordance with the laws of probability, and sometimes do not. This is the sort of view I shall advocate in the next section.[6]

[4] This is the view with which Ramsey contrasts his own view, in the discussion quoted earlier.

[5] I am imagining a case where due to the vagueness of the boundaries of horses, two horses are equally good candidates for having crossed the line first. In practice this would no doubt be deemed a tie, but imagine that we are examining very high-resolution pictures of the finish, and that we are interested not in the practical question of distributing winnings, but purely in the question of which horse in fact crossed the line first.

[6] Apart from my own view, another view which fits the description just given is that of Field (2000). Field supposes that an agent has a probability function P over propositions; he supposes also that the language includes a determinately operator D; and he then proposes that the agent's degree of belief $Q(\alpha)$ in any proposition α is given by $Q(\alpha) = P(D\alpha)$. Thus my degree of belief that α is my subjective probability that determinately α. It may sound, then, as though we *do* have

28.3 DEGREE OF BELIEF AS EXPECTED TRUTH VALUE

The picture I propose has three components: (1) an agent's epistemic state; (2) the degrees of truth of propositions; and (3) an agent's degrees of belief in propositions.

(1) I take an agent's epistemic state to be (represented by) a probability measure over the space of possible worlds. So, where W is the set of possible worlds, the agent's epistemic state P is a function which assigns real numbers between 0 and 1 inclusive to subsets of W. Intuitively, the measure assigned to a set S of worlds indicates how likely the agent thinks it is that the actual world is one of the worlds in S. Given this understanding of P—together with the convention that assigning a set of worlds measure 1 means that you are absolutely certain that the actual world is in that set, and assigning a set of worlds measure 0 means that you are absolutely certain that the actual world is not in that set—the three probability axioms are well motivated:

P1. For every set $A \subset W$, $P(A) \geq 0$

P2. $P(A \cup B) = P(A) + P(B)$ provided $A \cap B = \emptyset$

P3. $P(W) = 1$.

(2) At each possible world, each proposition has a particular degree of truth. Thus we may regard each proposition S as determining a function $S' : W \to [0, 1]$, i.e. the function which assigns to each world $w \in W$ the degree of truth of S at w.[7] The relationships between the functions associated with various propositions will be constrained in familiar ways by the logical relationships between these propositions: thus, for example, $(S \vee T)'(w) = \max\{S'(w), T'(w)\}$, $(S \wedge T)'(w) = \min\{S'(w), T'(w)\}$ and $(\neg S)'(w) = 1 - S'(w)$.

(3) We have a measure over worlds (the agent's epistemic state P), and functions from worlds to real numbers (each proposition S). Thus S is a *random variable*, and I propose that we identify the agent's degree of belief in S with her *expectation* (aka *expected value*) of S.

To get an intuitive feel for the proposal, consider the case where there are finitely many possible worlds. One's probability measure over sets of worlds is in this case determined, via the additivity axiom P2, by the values assigned to singleton sets: $P(\{w_1, \ldots, w_n\}) = P(\{w_1\}) + \ldots + P(\{w_n\})$. So, treating probabilities assigned to singletons as probabilities assigned to their members, one can, in the finite case, think

two different systems of degrees of belief: P-values and Q-values. But Field says that only Q-values are to be thought of as degrees of belief: 'P should be thought of as simply a fictitious auxiliary used for obtaining Q' (16); 'P [should] not be taken seriously: except where it coincides with Q, it plays no role in describing the idealized agent' (19). One worry I have about Field's proposal concerns the appearance of a primitive determinately operator within the contents of beliefs. A second worry concerns the downgrading of P: I think Field takes this too far. In my proposal (Section 28.3), subjective probabilities *do* play an important role in describing an agent, but they are *not* to be identified with degrees of belief. Field on the other hand seems to be in the grip of the view that if subjective probabilities are allowed into the picture at all (as anything beyond fictitious auxiliaries) then they will automatically grab the mantle 'degrees of belief'.

[7] For the sake of simplicity of presentation, I shall often conflate S and S', i.e. write of a proposition as *being* a function from worlds to degrees, rather than as *determining* such a function.

of oneself as assigning each *world* a degree of likelihood: a number indicating how likely one thinks it is that that world is the actual world. Each world w itself assigns each proposition S a degree of truth $S(w)$. Now, one's degree of belief in S is one's expectation of S, i.e. one's expected value of S's degree of truth. Let us denote this $E(S)$. In this finite case, it can be calculated as follows, where $w_1 \ldots w_n$ are all the possible worlds:

$$E(S) = P(\{w_1\}) \cdot S(w_1) + \ldots + P(\{w_n\}) \cdot S(w_n)$$

This is analogous to the calculation of expected utility in decision theory (with worlds playing the role of outcomes of acts, and degrees of truth playing the role of utilities of outcomes).

The proposal has two particularly important features: it meshes perfectly with the guiding idea of one's degree of belief that S as a measure of the strength of one's tendency to act as if S; and it has the consequence that degrees of belief sometimes behave like probability assignments, and sometimes do not. I shall discuss these points in turn.

First, consider the idea that one's degree of belief that S is a measure of the strength of one's tendency to act as if S. It is important to note that I am not claiming that two persons who have the same degree of belief that S will *behave in the same ways*, or even have the same *tendencies* to behave in certain ways. I am claiming that they will have the same tendency to *act as if S*. Whether a person's behaving in a certain way constitutes her *acting as if S* depends on her preferences (desires, utilities) and on her other beliefs. For example, let S be the proposition that there is an especially fragrant rose in Bob's garden. For a rose fancier, approaching Bob's garden might constitute acting as if S, whereas for a person with an aversion to roses—or a rose fancier with false beliefs about the location of Bob's garden—moving away from Bob's garden might constitute acting as if S. So while two persons who have the same degree of belief that S will have the same tendency to act as if S—this is our guiding idea—in general they will only *behave in the same ways* (described at the level of bodily movements, for example—rather than in terms of whether they are acting as if S) if their *other* beliefs and desires are *also* the same.

Consider now a simple example. There are three 'open worlds' w_1, w_2 and w_3—i.e. three worlds such that one is *not* certain that one is *not* in them—i.e. $P(\{w_1, w_2, w_3\}) = 1$. Suppose that S is the proposition 'A tall person will win the race'. You don't know who will win, but you do know that it is either the first man in our Sorites series leading from tall men to short men (this is the situation in w_1), or the last man (this is the situation in w_2), or the man in the middle (this is the situation in w_3). You think that each of these three possibilities is equally likely, i.e. $P(\{w_1\}) = P(\{w_2\}) = P(\{w_3\}) = \frac{1}{3}$. In w_1, S is 1 true; in w_2, S is 0 true; in w_3, S is 0.5 true. So your expectation that S is $\frac{1}{3} \cdot 1 + \frac{1}{3} \cdot 0 + \frac{1}{3} \cdot 0.5 = 0.5$. This seems to be a true measure of the strength of your tendency to act as if S. Suppose you need a tall man for your basketball team, and you have a choice between signing up the race winner (whoever that should turn out to be), or Bill (whom you know to be of the same height as the first man in our Sorites series—hence 'Bill is tall' is 1 true, and you know this, and so your expectation of this proposition is 1), or Ben (whom you know to be of the same height as the last man in our Sorites series—hence 'Ben is tall' is 0 true, and you know this, and so your expectation of this proposition is 0), or Bob (whom you know to be of the same height as the man in the middle of our Sorites series—hence 'Bob is tall' is 0.5 true, and you know this, and so your expectation of this proposition is 0.5). You would sooner sign up the race winner than Ben, sooner sign up Bill than the race winner, and be indifferent between signing up the race winner and Bob. Thus, the strength of your tendency to act as if *S* mirrors your expectation of *S*.

I have been making the assumption that your preferences regarding team members can be summed up thus: 'The taller the better.' If, on the other hand, you wanted only *very* tall players—so you are just as averse to signing up a borderline tall person as to signing up a short person—then of course you would have *no* tendency to sign up Bob. That is no problem for my view (even though, in this new case—in which you have different preferences—your expectation that Bob is tall is still 0.5). For if you wanted only *very* tall players, then signing up *P* would *not* constitute acting as if *P* is tall; rather, it would constitute acting as if *P* is *very* tall (recall the discussion on p. 497). In the situation described (in both cases—i.e. whatever your preferences), your expectation that Bob is *very* tall is 0. So in the second case—where (given your new preferences) signing up *P* now constitutes acting as if *P* is very tall, rather than acting as if *P* is tall—my theory *correctly* predicts that you will have *no* tendency to sign up Bob. In sum: as your preferences change from 'the taller the better' to 'very tall', your degrees of belief that Bob is tall and that Bob is very tall remain 0.5 and 0 respectively. However, the significance of signing-up behaviour changes. At first, such behaviour constitutes acting as if the signed-up player is tall; with the new preferences, it constitutes acting as if the signed-up player is very tall. That is why two people who have the *same* degree of belief in 'Bob is tall' might have *different* tendencies to *sign up Bob*. My claim is that they will have the *same* tendency to *act as if Bob is tall*. If their preferences differ, however, then what counts as acting as if Bob is tall for one person—say, signing up Bob—might not count as acting as if Bob is tall for the other person.

The same kind of point applies in a host of other cases which, at first sight, might seem to pose a problem for my view. For example, suppose that persons *A* and *B* are faced with a choice of cups of coffee: cup 1, which they know is either freshly made or has been sitting there for several hours (they do not know which—but they do know each option is equally likely), or cup 2, which they know was made about fifteen minutes ago. In the circumstances, we may suppose that both *A* and *B* assign an expected truth value of 0.5 to both 'cup 1 is hot' and 'cup 2 is hot'—i.e. on my view both *A*'s and *B*'s degrees of belief in both these propositions are 0.5. But *A* and *B* behave quite differently. *A*, who likes her coffee either very hot, or cooled to room temperature, reaches for cup 1 and has no tendency whatsoever to reach for cup 2. *B*, whose preference in coffee is 'the hotter the better', is equally inclined to reach for cup 1 as for cup 2. This is all grist for my mill. Both *A* and *B* believe to degree 0.5 that the coffee in cup 2 is hot, and believe to degree 0 that the coffee in cup 2 is very hot. Given *B*'s preferences, *reaching for a cup* is (other things being equal) a way of *acting as if it contains hot coffee*. Given *A*'s preferences, *reaching for a cup* is *not* (other things being equal) a way of *acting as if it contains hot coffee*; rather, it is (other things being equal) a way of *acting as if it contains coffee which is very hot or at room temperature*. *A* and *B* have the same degree of belief that cup 2 contains hot coffee. So my claim is that they will have the *same* tendency to *act as if cup 2 contains hot coffee*. But 'acting as if a cup contains hot coffee' amounts to doing something different in *A*'s case than in *B*'s. That is why *A* and *B* have different tendencies to *reach for cup 2*, even though they have the same degree of belief that it contains hot coffee, and the same tendency to act as if it contains hot coffee. The key point, then, is the one made on p. 497: whether a person's *behaving in a certain way* constitutes her *acting as if S* depends on her preferences and on her other beliefs. So two people who have the same degree of belief that *S*, but differ in their other beliefs or in their preferences, might behave differently (described at the level of bodily movements), even though they have the same tendency to act as if *S*.

Apart from meshing with the idea of one's degree of belief that *S* as a measure of the strength of one's tendency to act as if *S*, my proposal also has the desired feature that sometimes degrees of belief behave like probability assignments, and sometimes do not. Before showing this, I shall generalize the picture presented above. For so far we have considered only the special case where we have finitely many possible worlds, but of course we cannot, in general, suppose that there are only finitely many possible worlds—indeed we cannot suppose that there are only countably many. But if there are uncountably many possible worlds, then (i) we cannot assume that the agent's probability measure is defined on all subsets of the space of possible worlds, and (ii) we cannot assume that every proposition determines a *measurable* function from worlds to truth values, i.e. a random variable. We shall handle this situation in the standard way. In regards to point (i), we suppose there to be a family $\mathcal{F}$ of subsets of the space W of all possible worlds which is a σ-field, i.e. it satisfies the conditions:

1. $W \in \mathcal{F}$
2. For all $A \in \mathcal{F}, \overline{A} \in \mathcal{F}$
3. For any countable number of sets $A_1, \ldots, A_n$ in $\mathcal{F}$, $\bigcup_n A_n \in \mathcal{F}$.[8]

Our probability measure will be defined on $\mathcal{F}$, i.e. it will assign probabilities to sets in $\mathcal{F}$, and not to other subsets of W; the sets in $\mathcal{F}$ will be called the measurable sets of possible worlds.[9] In regards to point (ii), for a function S from worlds to the reals to be measurable, i.e. a random variable, it must satisfy the condition that for any real x, $\{w \in W : S(w) \leq x\} \in \mathcal{F}$. If such a function is bounded, it will have a well-defined expectation $E(S)$. All propositions are functions from worlds to $[0, 1]$, and hence bounded. As for the condition that they be measurable, we henceforth restrict our attention to propositions which meet it. This means that we consider only propositions *S* such that it makes sense to ask 'How likely do you take it to be that this proposition has a truth value within such-and-such limits?'

[8] By de Morgan's laws, we could equivalently replace union with intersection in condition 3.

[9] Once we have made this alteration to our set-up, it is standard also to change axiom P2 so that it applies not just to unions of two sets, but to unions of countably many sets—i.e. for any countable collection $\{A_i\}$ of pairwise disjoint sets, $P(\bigcup_{n=1}^{\infty} A_n) = \sum_{n=1}^{\infty} P(A_n)$.

With the general picture now in place, we can make the following definitions:

Definition 1 (vagueness-free situation). An agent is in a vagueness-free situation (VFS) with respect to a proposition S iff there is a measure-1 set T of worlds (i.e. a set T such that $P(T) = 1$) such that $S(w) = 1$ or $S(w) = 0$ for every $w \in T$. (That is, the agent may not know for sure whether S is true or false, but she does absolutely rule out the possibility that S has an intermediate degree of truth: for she is certain that the actual world is somewhere in the class T, and everywhere in T, S is either 1 true or 0 true.) An agent is in a VFS with respect to a *set* Γ of propositions if she is an a VFS with respect to each of the propositions in Γ.

Definition 2 (uncertainty-free situation). An agent is in an uncertainty-free situation (UFS) with respect to a proposition S iff there is a measure-1 set T of worlds and a $k \in [0, 1]$ such that $S(w) = k$ for every $w \in T$. (That is, it is totally ruled out that S has a degree of truth other than k: for the agent is certain that the actual world is somewhere in the class T, and everywhere in T, S is k true.) An agent is in a UFS with respect to a *set* Γ of propositions if she is an a UFS with respect to each of the propositions in Γ.

We can now establish four results which show when degrees of belief behave like probability assignments, and when they do not.

Proposition 1 (Degrees of belief equal probabilities in VFSs). If an agent is in a VFS with respect to S, then $E(S) = P(\{w : S(w) = 1\})$.[10]

Proposition 2 (Degrees of belief equal degrees of truth in UFSs). If an agent is in a UFS with respect to S, then $E(S)$ equals the degree of truth which the agent is certain S has.[11]

Proposition 3 (Degrees of belief behave like probabilities in VFSs). Let Γ be a class of wfs, closed under the operations of forming wfs using our standard propositional connectives $\vee$, $\wedge$ and $\neg$, such that one is in a VFS with respect to Γ.[12] Then one's

[10] *Proof.* We are given that there is a set T of worlds such that $P(T) = 1$ and $S(w) = 1$ or $S(w) = 0$ for every $w \in T$. Divide T into two sets: T_1, containing the worlds in which S is 1 true, and T_0, containing the worlds in which S is 0 true. (We know these are both measurable as follows. Where S is a random variable and a is any real, $\{w : S(w) = a\}$ is measurable. $T_0 = T \cap \{w : S(w) = 0\}$, and $T_1 = T \cap \{w : S(w) = 1\}$, and measurable sets are closed under intersection.) The expectation of a formula is not affected by its truth value anywhere outside a measure 1 set, so $E(S) = P(T_0) \cdot 0 + P(T_1) \cdot 1 = P(T_1)$. Let S_1 be $\{w : S(w) = 1\}$. $P(S_1) = P(T_1) + P(S_1 \backslash T_1)$. But $P(S_1 \backslash T_1) = 0$, because $P(T) = 1$ and so if the measure of some set disjoint from T were positive, then by P2 the measure of W would be greater than 1, violating P3. So $P(S_1) = P(T_1)$ and hence $E(S) = P(S_1)$.

[11] *Proof.* There is a measure 1 set T such that $S(w) = k$ for every $w \in T$. The expectation of a formula is not affected by its truth value anywhere outside a measure 1 set, so $E(S) = P(T) \cdot k = 1 \cdot k = k$.

[12] The closure requirement is no restriction, because if one is in a VFS with respect to a class of wfs, then one is in a VFS with respect to the closure of that class (because whenever the component wfs are 1 true or 0 true, so are the compounds).

degrees of belief (i.e. expectations) of wfs in Γ behave like probabilities, in the sense that they satisfy the following three conditions:

1. For all wfs $\gamma \in \Gamma$, $0 \leq E(\gamma) \leq 1$.[13]
2. For all tautologies $\gamma \in \Gamma$, $E(\gamma) = 1$.[14]
3. If $\gamma_1, \gamma_2 \in \Gamma$ are mutually exclusive, then $E(\gamma_1 \vee \gamma_2) = E(\gamma_1) + E(\gamma_2)$.[15]

Proposition 4 (Degrees of belief behave like degrees of truth in UFSs). Let Γ be a class of wfs, closed under the operations of forming wfs using $\vee$, $\wedge$ and $\neg$, such that one is in a UFS with respect to Γ.[16] Then one's degrees of belief (i.e. expectations) of wfs in Γ behave like degrees of truth, in the sense that they satisfy the following three conditions:

1. $E(\neg\gamma) = 1 - E(\gamma)$.[17]
2. $E(\gamma_1 \vee \gamma_2) = \max\{E(\gamma_1), E(\gamma_2)\}$.
3. $E(\gamma_1 \wedge \gamma_2) = \min\{E(\gamma_1), E(\gamma_2)\}$.[18]

[13] *Proof.* By proposition 1, $E(\gamma) = P(\{w : \gamma(w) = 1\}$. As this is a probability, it is between 0 and 1 (inclusive) by definition.

[14] There are several possible definitions of 'tautology' in fuzzy logic. All we need for the proof is something they all agree on, viz. that a tautology never gets the value 0. *Proof.* By hypothesis we have a set T of worlds such that $P(T) = 1$ and $\gamma(w) = 1$ or $\gamma(w) = 0$ for every $w \in T$. But as γ is a tautology, there are no worlds w such that $\gamma(w) = 0$, so we have a set T of worlds such that $P(T) = 1$ and $\gamma(w) = 1$ for every $w \in T$. So $E(\gamma) = 1$.

[15] There are several possible definitions of 'mutually exclusive' in fuzzy logic. All we need for the proof is something they all agree on, viz. that two mutually exclusive propositions never both get the value 1. *Proof.* By hypothesis we have a set T_1 of worlds such that $P(T_1) = 1$ and $\gamma_1(w) = 1$ or $\gamma_1(w) = 0$ for every $w \in T_1$, and a set T_2 of worlds such that $P(T_2) = 1$ and $\gamma_2(w) = 1$ or $\gamma_2(w) = 0$ for every $w \in T_2$. So $P(T_1 \cap T_2) = 1$. (For suppose it has measure $0 \leq n < 1$. Then $T_1 \backslash T_2$ and $T_2 \backslash T_1$ both have measure $1 - n$. But then by P2, $(T_1 \backslash T_2) \cup (T_2 \backslash T_1) \cup (T_1 \cap T_2)$ has measure $(1 - n) + (1 - n) + n = 2 - n > 1$, for these three sets are pairwise disjoint.) So we have a measure 1 set $T_1 \cap T_2$ in which both γ_1 and γ_2 are 0 or 1 true at *every* world. But we also know γ_1 and γ_2 are mutually exclusive, i.e. there are no worlds where γ_1 and γ_2 are both 1 true. So we can divide our measure 1 set $T_1 \cap T_2$ into three pairwise disjoint subsets, G, G_1 and G_2, with G containing worlds at which both γ_1 and γ_2 are 0 true, G_1 containing worlds at which γ_1 is 1 true and γ_2 is 0 true, and G_2 containing worlds at which γ_2 is 1 true and γ_1 is 0 true. (We know these subsets are measurable by reasoning similar to that in n.10. Note also that if the set of worlds where each atomic formula is true is measurable, then by the conditions on a σ-field, the set of worlds where each propositional compound is true is also measurable.) The expectation of a formula is not affected by its truth value anywhere outside our measure 1 set $T_1 \cap T_2 (= G \cup G_1 \cup G_2)$. So $E(\gamma_1) = P(G_1)$, $E(\gamma_2) = P(G_2)$, and $E(\gamma_1 \vee \gamma_2) = P(G_1) + P(G_2)$ (because $\gamma_1 \vee \gamma_2$ is true at worlds in G_1 and G_2 and false at worlds in G). Hence $E(\gamma_1 \vee \gamma_2) = E(\gamma_1) + E(\gamma_2)$.

[16] Again, the closure requirement is no restriction, because if one is in a UFS with respect to a class of wfs, then one is in a UFS with respect to the closure of that class (because if one is certain that S is m true and that T is n true, then one is certain that $S \vee T$ is $\max\{m, n\}$ true, that $S \wedge T$ is $\min\{m, n\}$ true, and that $\neg S$ is $1 - m$ true).

[17] *Proof.* There is a measure 1 set T such that at every world in T, γ is k true. So $E(\gamma) = k$. At every world in T, $\neg\gamma$ is $1 - k$ true. So $E(\neg\gamma) = 1 - k = 1 - E(\gamma)$.

[18] *Proofs.* By hypothesis we have a set T_1 of worlds such that $P(T_1) = 1$ and $\gamma_1(w) = m$ for every $w \in T_1$, and a set T_2 of worlds such that $P(T_2) = 1$ and $\gamma_2(w) = n$ for every $w \in T_2$. So $P(T_1 \cap T_2) = 1$, as in n.15. At every world in $T_1 \cap T_2$, γ_1 is m true and γ_2 is n true, hence $\gamma_1 \vee \gamma_2$ is $\max\{m, n\}$ true and $\gamma_1 \wedge \gamma_2$ is $\min\{m, n\}$ true. The expectation of a formula is not affected by its

Summing up my proposal: an agent's degrees of belief are the resultant of two things: the agent's uncertainty about which way the actual world is (represented by a probability measure over the space of possible worlds, with the measure assigned to a set of worlds specifying how likely the agent thinks it is that the actual world is in that set), and the facts about how true each proposition is in each world. Specifically, the agent's degree of belief in a proposition is the agent's expected value of its degree of truth: roughly, the average of its truth in all the worlds the agent has not ruled out, weighted according to how likely the agent thinks it is that each of those worlds is the actual one. In some situations, the agent will have ruled out vagueness: she may not know which world is actual, but she is certain that in the actual world, some propositions of interest are either fully true or fully false. In such situations, her degrees of belief will behave like probabilities (propositions 1 and 3). In other situations, the agent will be free of uncertainty with respect to some propositions of interest: she is certain of exactly how true they are in the actual world. In such situations, her degrees of belief will behave like degrees of truth (propositions 2 and 4). In situations which are neither vagueness-free nor uncertainty-free—that is, where the agent is unsure of the truth values of some propositions of interest, and cannot rule out vagueness, that is, cannot rule out that they might have intermediate degrees of truth—her degrees of belief in those propositions need not behave like probabilities or degrees of truth. (In situations which are *both* uncertainty-free *and* vagueness-free—that is, the agent knows of each of the propositions in question that it is 1 true, or that it is 0 true—degrees of belief behave both like probabilities and like degrees of truth. This is possible because the behaviours of probabilities and degrees of truth coincide in this special case.) In *all* cases, I maintain that an agent's expectation of a proposition S's degree of truth is an accurate measure of her tendency to act as if S, and this is why I identify degrees of belief with expectations.

My proposal contrasts with the standard view, as expressed for example in the following passages:

> Let our degrees of belief be represented by a probability measure, P, on a standard Borel space (Ω, F, P), where Ω is a set, F is a sigma-field of measurable subsets of Ω, and P is a probability measure on F.
>
> (Skyrms, 1984, 53)

> [By a reasonable initial credence function C] I meant, in part, that C was to be a probability distribution over (at least) the space whose points are possible worlds and whose regions (sets of worlds) are propositions. C is a non-negative, normalized, finitely additive measure defined on all propositions.
>
> (Lewis, 1986, 87–8)

The crucial difference between the standard view and mine is that the former equates an agent's degrees of belief directly with her subjective probabilities. My view, on the other hand, countenances the subjective probability measure—it models the agent's epistemic state—but regards degrees of belief as resultants of this state and the truth

truth value anywhere outside our measure 1 set $T_1 \cap T_2$, so $E(\gamma_1) = m$, $E(\gamma_2) = n$, $E(\gamma_1 \vee \gamma_2) = \max\{m, n\} = \max\{E(\gamma_1), E(\gamma_2)\}$, and $E(\gamma_1 \wedge \gamma_2) = \min\{m, n\} = \min\{E(\gamma_1), E(\gamma_2)\}$.

values of propositions at worlds. In the sort of cases Skyrms and Lewis were considering, in which bivalence was assumed, this difference makes no difference (propositions 1 and 3). However, if we want to add vagueness to the mix, then we will run into all sorts of problems, if we have already identified degrees of belief with subjective probabilities—for, as we saw at the outset, vagueness also gives rise to degrees of belief, but these degrees of belief do not behave like probabilities. On the other hand, if we identify degree of belief with expected truth value even in the bivalent case, then we can generalize smoothly to the case of vagueness, handled using degrees of truth.

28.4 OBJECTIONS AND REPLIES

(1) If your degrees of belief do not conform to the probability calculus, then you are subject to Dutch book, i.e. you are irrational. Reply: One should not bet at all on a proposition S unless one is in a vagueness-free situation with respect to S; if one does bet in a non-VFS, then it is for that reason alone that one is irrational. Suppose you are not in a VFS with respect to S. Suppose first that you know that S is k true, for some $k \in (0, 1)$; say $k = 0.5$ for the sake of argument. Then you should not bet on S. For to bet is to agree to an arrangement whereby you get such-and-such *if S turns out to be the case*. But you already know what is the case—and you know that it is, in the nature of things, indeterminate whether S—hence indeterminate whether you get your payoff. Knowing all this, you should not bet in the first place. Second, suppose that you do not know whether S is true—and you cannot rule out that S has an intermediate degree of truth. In this case again you should not bet, because for all you know, the bet will not—for the sort of reason just seen—be able to be decided. Of course if there is in place some system for deciding bets on S when S has an intermediate degree of truth—say an umpire who rules one way or the other, or a rule that S will be deemed 1 true if it is more than 0.5 true—*then* one may enter into a betting arrangement on S. However in such a case the situation has, in effect, been turned into a VFS, by changing S's intermediate degrees of truth in some non-ruled-out worlds into 1's or 0's.[19]

[19] My comments about not betting in non-VFSs are concerned with standard bets—i.e. bets which do not specify what is to happen (who gets what) when the proposition in question is neither true nor false. Milne (2007) discusses a new type of betting arrangement, tailor-made for vagueness, on which one *could* legitimately bet in a non-VFS. The basic idea (although this is not the way Milne expresses it) is that if one bets on S, and S is n true, then one receives n times the stake. Of course this complements rather than conflicts with my comments above (Milne was not suggesting otherwise). I say that one should not accept an *ordinary* bet if one thinks that vagueness may be present—for when vagueness is involved, there is no way of deciding such a bet. This does not mean that one should not accept a new kind of bet—one designed precisely to avoid the problem faced by ordinary bets when vagueness is present, by explicitly building in a decision procedure which works even when the proposition on which one is betting has an intermediate degree of truth. (When I was writing the paper on which Milne (2007) was a comment, I considered the idea of introducing a type of bet along the lines discussed by Milne, designed specifically to handle vague outcomes. I did not pursue the idea, however, because in general we do not know the precise degrees of truth of vague propositions, so even if it is fixed that when S is n true, one receives n times the stake, still we will in general have no way of actually deciding and paying out the bet, because we will not know

(2) Some writers have claimed that 'The cunning bettor is simply a dramatic device—the Dutch book a striking corollary—to emphasize the underlying issue of coherence' (Skyrms, 1984, 22). The idea is meant to be that one is internally incoherent if one's degrees of belief do not conform to the probability calculus: the Dutch book idea simply serves to bring this incoherence into the open in a striking way; but even if one is not subject to Dutch book for some reason (e.g. because betting has been made illegal and this law is enforced absolutely) one is still internally incoherent. Reply: Why is one supposed to be incoherent in such a case? Well, here's a way of bringing it out. Suppose I think *A* is 50% likely to occur (in 50% of futures compatible with the present, *A* occurs); I think *B* is 50% likely to occur (in 50% of futures compatible with the present, *B* occurs); I think *A* and *B* are incompatible (in no future do *A* and *B* both occur); and yet I think '*A* or *B*' is not 100% likely to occur—i.e. I think that in (say) 50%, rather than 100%, of futures compatible with the present, '*A* or *B*' will be true. When framed in this way in terms of sizes of sets of possible futures, this combination of beliefs is obviously incoherent. But my view endorses this assessment: in the situation envisaged, the agent is in a VFS (she does not know whether or not *A* or *B* will occur, but she assumes neither of them will sort-of occur), and so will not have these degrees of belief, on my view. On the other hand I do *not* think that, in itself, the following combination of degrees of belief is incoherent, even supposing the agent knows that *A* and *B* cannot both be fully true:

$$A : 0.5, \quad B : 0.5, \quad A \text{ or } B : 0.5$$

It all depends on how these degrees of belief arise. If you are in a VFS and have these degrees of belief, then you are indeed incoherent—as can be brought out either by Dutch book reasoning, or by reflections on sizes of sets of possibilities. But degrees of belief might arise in other ways—not just as a result of uncertainty; and when they do, this sort of combination can be perfectly reasonable. For example, suppose that *A* is the proposition that a certain leaf is red, and *B* is the proposition that it is orange; then *A* and *B* cannot both be fully true. Suppose also that the leaf in question is right in the middle of a Sorites series leading from red things to orange things. Then, I submit, the above combination of degrees of belief is perfectly reasonable: intuitively it is just fine, and neither the Dutch book nor the 'sizes of sets of possibilities' rationales can get a grip to show that there is something wrong with it. Dutch book reasoning does not get started because I will not bet (there is nothing to bet on—no outcome to wait and see about: I already have all the information about the leaf's colour before me). Similarly, the 'sizes of sets of possibilities' reasoning does not get started, because there is nothing I am uncertain about.

(3) I claim that in non-VFSs, we have degrees of belief while not being prepared to bet (at all). The objection is that we cannot make sense of the idea of degrees of belief except in terms of fair betting quotients or odds. Reply: We make sense of the idea of degree of belief in *S* in terms of strength of tendency to act as if *S*, and 'acting

the actual value of *n* in question. However Peter pointed out to me that we are often in a similar position with regards to ordinary bets—i.e we cannot determine the outcome—but this does not reduce their theoretical interest in relation to degrees of belief.)

as if *S*' can be made sense of more generally than in terms of 'betting on *S*'. After all, betting is essentially tied up with uncertainty—betting gets its life from the fact that we do not know what the outcome will be—but, I have argued, the idea of degree of belief gets a grip in circumstances in which there is no uncertainty at all. Consider an autumn leaf which is borderline red–orange. You have some tendency to act as if it is red, as discussed below (objection 5). But with the leaf in plain sight, you would not accept a bet that it is red, at any price: for we can all see quite plainly that the leaf is neither clearly red nor clearly non-red, and so we can see at the outset that the bet will misfire.[20]

(4) Suppose my degree of belief in *S* is 0.5 because I am uncertain whether *S* is 1 true or 0 true. Then I might buy a bet on *S*, if the price and prize are right. But suppose my degree of belief in *S* is 0.5 because I am certain that *S* is 0.5 true. Then, for the reasons discussed above, I will not buy a bet on *S*, no matter what the price or prize. So the same state—a degree of belief of 0.5 in *S*—leads to different actions. How can this be, if my degree of belief measures my tendency to act as if *S*? Reply: These different actions are the results not of a single belief, but of complexes of beliefs, which are different in the two situations. A 0.5-degree belief that *S* combined with the belief that whatever further evidence comes in, I will not alter my degree of belief in *S*, leads to refusing to bet; a 0.5-degree belief that *S* combined with the belief that further evidence might come in leading me to believe to degree 1 that *S*, and that further evidence might come in leading me to believe to degree 0 that *S*, leads to accepting certain bets.

(5) One's expectation that *S* is not an accurate measure of one's tendency to behave as if *S*. Suppose I know that a certain orangey-red autumn leaf is red to degree 0.5. Suppose also that I need a perfectly red leaf. Then I will have no tendency whatsoever to reach for this leaf, even though my expectation that it is red is 0.5. Reply: The key here is the presence of the word 'perfectly'. Of course if I need a *perfectly* red leaf, then I will have no tendency whatsoever to reach for the orangey-red one. But this is quite compatible with the foregoing account, because my expectation that the leaf is perfectly red, i.e. red to degree 1, is 0. On the other hand, my expectation that it is red is 0.5; and if I need a *red* leaf, then I think I would have some tendency to reach for this one: less than for a perfectly red leaf, but more than for a green one.[21]

References

Field, Hartry (2000), 'Indeterminacy, degree of belief, and excluded middle', *Noûs* 34, 1–30.

Lewis, David (1986), 'A subjectivist's guide to objective chance' in *Philosophical Papers*, vol. II, 83–132, Oxford University Press, New York.

MacFarlane, John (2006), 'The things we (sorta kinda) believe', *Philosophy and Phenomenological Research* 73, 218–24.

[20] Those who feel strongly that where there are degrees of belief there must be betting quotients can find comfort in the kind of betting arrangement discussed in Milne (2007) (see n. 19 above). Milne shows that the fair betting quotient a rational agent assigns to a bet on *A* of his kind perfectly matches the agent's degree of belief that *A* in my sense, i.e. her expectation of *A*'s degree of truth.

[21] Further objections to my view are considered in Smith (2008, §5.3.3).

Milne, Peter (2007), 'Bets and fuzzy propositions: Comments on Nicholas J. J. Smith's "Degrees of truth, degrees of belief, and pragmatics",' presented at the Arché Vagueness Conference, St Andrews, 8 June.

Ramsey, F. P. (1990), 'Truth and probability' in D. H. Mellor, ed., *Philosophical Papers*, 52–94, Cambridge University Press, Cambridge.

Schiffer, Stephen (2000), 'Vagueness and partial belief,' *Philosophical Issues* 10, 220–57.

Skyrms, Brian (1984), *Pragmatics and Empiricism*, Yale University Press, New Haven.

Smith, Nicholas J.J. (2008), *Vagueness and Degrees of Truth*, Oxford University Press, Oxford.