

Belief, Cognition and the Will

Anthonie Meijers (ed.)



Tilburg University Press 1999

Luc Bovens

1 INTRODUCTION

I can express my doxastic attitudes with respect to some proposition P both in qualitative and in quantitative terms. In qualitative terms, I might say that I believe that P, do not believe that P or believe that not-P. In quantitative terms, I might say that I ascribe a certain degree of confidence to P, i.e. that I assign a subjective probability to P. Is there a systematic relationship between these two types of expressions of doxastic attitudes? Since our quantitative vocabulary is more fine-grained than our qualitative vocabulary, one could conjecture that a relationship of supervenience holds. The supervenience thesis could be roughly stated as follows: if two persons have the same quantitative doxastic attitudes, then they must also have the same qualitative doxastic attitudes. My goal is to investigate whether the supervenience thesis can be upheld or breaks down for particular interpretations of the relationship between qualitative and quantitative doxastic attitudes.

The connection between beliefs and degrees of confidence is by no means equivocal. Without trying to be exhaustive, I propose to provide a number of interpretations of *belief* which bring out the connection between our qualitative and quantitative talk about doxastic attitudes:

(i) First, there is the *dogmatic* notion of belief. To believe that P is to assign a subjective probability of 1 to P. We have a special qualitative term for this doxastic attitude as well, viz. to be *certain* that P. On the dogmatic view, to believe that P is simply to be certain that P.

(ii) Second, there is the *Lockean* notion of belief. Believing that P is to assign a sufficiently high subjective probability to P, which may however fall short of 1. This probability should be somewhere in the neighborhood of 1, but the threshold value may vary from one person to another for a particular proposition P. Presumably, we should not extend the neighborhood of 1 to values that are smaller than or equal to .50, since then we would be believing both P and not-P. Hence, to believe that P entails that one at least takes P more likely to be so than not.

(iii) Third, there is the *agentic* notion of belief. Upon being asked whether I believe that there is a witch dwelling in the woods, I might respond: "I believe enough not to go up there."¹ This doxastic attitude may be entirely consistent with believing it more likely that there is no witch dwelling in the woods than that there is a witch dwelling in the woods. Nonetheless, I am not willing to go there, since why would I want to take any chances? Since my actions are indistinguishable from someone who is convinced that there is a witch living in the woods (and, like me, fears witchcraft), I am said to believe just the same that there is a witch living in the woods.

(iv) Fourth, there is the *abductive* notion of belief. Consider some area of science that experiences lots of change. For example, suppose that some part of the phylogenetic tree is radically being redrawn with every new paleontological find that is relevant to the problem at hand. Then I might say that I believe the newest version, because of all the possible drawings of this part of the tree, this one is most likely to be true, since it provides the best explanation of the data. But this does not mean that I think it more likely than not that this drawing in all of its details is correct: It may well be the case that I consider the disjunction of all the alternative tree-drawings more plausible than this newest version.

My claim is that to make sense of our belief talk in ordinary life we need to have a range of notions at our disposal. I am not claiming that we may freely shift between these different notions of belief. These notions have their use within particular contexts of discourse and it may raise some eyebrows if we were to use them outside of these contexts. It would be odd to say that I do not believe that the airliner that I am about to board will arrive safely and yet the dogmatic view of belief could account for such an utterance. It would be odd to say that I believe that I can draw a non-face card at random from a regular deck of cards and yet the Lockean view of belief could account for such an utterance. It would be odd to say that I believe that I will get sick during my trip abroad when I fill out the insurance papers at the travel agency, and yet the agentic notion of belief could account for such an utterance. And finally, it would be odd to say that I believe that today's winner in the casino played blackjack rather than the roulette without having any more information at my disposal that the odds in blackjack are somewhat better than the odds in roulette, and yet the abductive notion of belief could account for such an utterance. The following scope distinction clarifies my position. For all contexts *C*, there is some notion of belief *N* such that *N* makes sense of our ordinary belief talk in *C*. But it is not true that there is some notion of belief *N*, such that in all contexts *C*, *N* can make sense of our ordinary belief talk in *C*.

Let us now turn to the supervenience thesis. A propositional variable *R* (written in *italic script*) can take on two values, viz. *R* and not-*R* (written in roman script). Let there be some set of propositional variables $S = \{R_1, \dots, R_n\}$. To say that two people are in the same quantitative doxastic state with respect to *S* is

¹ The line is taken from the movie "The Blair Witch Project".

to say that they share one and the same (subjective) probability distribution over these propositional variables. To say that two people are in the same qualitative doxastic state with respect to *S* is to say that their belief sets include precisely the same propositions that are the values of the propositional variables in *S*. The supervenience thesis states that, necessarily, if two people are in the same quantitative state with respect to *S*, then they are in the same qualitative state with respect to *S*. It turns out that the supervenience thesis is true for some, but not for other notions of belief. Our question is: if the supervenience thesis is false for some notion of belief, can we determine what other properties should be included in the base beside the quantitative doxastic states in order to make the supervenience thesis come out true?

The answer to this question has repercussions on the debate between evidentialism and doxastic voluntarism. Should our beliefs be strictly determined by the available evidence, or are there certain notions of belief in which valuational attitudes have a legitimate role to play? If the latter is true, what sort of attitudes could take up this role? How does this affect the question whether one may *will* to believe? I conclude with a speculation about what doxastic supervenience theses do to the status of beliefs *versus* degrees of confidence and where this places me in the debate between evidentialists and voluntarists.

2 THE DOGMATIC NOTION OF BELIEF

For all values *R* of the variables in *S*, I dogmatically believe that *R* just in case the marginal subjective probability of *R* is 1. It is easy to see that the supervenience thesis holds on the dogmatic view of belief, since the probability distribution over a set of variables uniquely determines the content of the belief set. Dogmatic believers believe all and only the propositions whose marginal probability is 1 and remain agnostic about the remaining ones. The supervenience thesis makes this an attractive notion, but the drawback is that we would come to believe very few things in life.

3 THE LOCKEAN² NOTION OF BELIEF

This notion is more liberal. For all values *R* of the variables in *S*, I hold a Lockean belief that *R* just in case the marginal subjective probability of *R* is greater than or equal to some threshold value. The supervenience thesis does not hold for this notion of belief, since two persons may share the same probability distribution and yet have different beliefs, because they do not share the same threshold value. What is it that determines this threshold value for belief? For

² The connection to Locke was proposed by Richard Foley (1992). Hawthorne and Bovens (1999) present an exploration of the precise relationship between the Lockean notion of belief and degrees of confidence.

some proposition R, I may either remain agnostic with respect to R or I may come to believe that R, and in the latter case, either my belief will be true or my belief will be false. We construct a utility function over the states of the world in which I truly believe that R, remain agnostic with respect to R and I falsely believe that R. Epistemic integrity requires that I prefer a state of affairs in which I truly believe that R to a state of affairs in which I am agnostic with respect to R to a state of affairs in which I falsely believe that R.³ But it does not tell us anything about the strength of these preferences: the strength may differ depending on the person and depending on whether we are talking about scientific beliefs, religious beliefs... We can define the following Von Neumann-Morgenstern utility function⁴ over these states of affairs:

$$\begin{aligned} U(\text{True-Belief}) &= 1 \\ U(\text{Agnosticism}) &= a \text{ for } 0 < a < 1 \\ U(\text{False-Belief}) &= 0 \end{aligned}$$

If I am just about indifferent between True-Belief and Agnosticism but strongly prefer Agnosticism to False-Belief, then a will be close to 1. If I strongly prefer True-Belief to Agnosticism, but am just about indifferent between Agnosticism and False-Belief, then a will be close to 0. We can now calculate the expected utility of *believing* that R and of *refraining from believing that R and from believing that not-R*:

$$\begin{aligned} EU(\text{believing-that-R}) &= U(\text{True-Belief})P(R) + U(\text{False-Belief})P(\text{not-R}) = P(R) \\ EU(\text{refraining-from-...}) &= U(\text{Agnosticism}) \\ EU(\text{believing-that-not-R}) &= U(\text{True-Belief})P(\text{not-R}) + U(\text{False-Belief})P(R) = \\ &P(\text{not-R}) \end{aligned}$$

If $U(\text{Agnosticism}) < .50$, then we adopt either the belief that R or the belief that not-R, since either $P(R) \geq .50$ or $P(\text{not-R}) \geq .50$. But as $U(\text{Agnosticism})$ rises, we are less prone to adopt beliefs: we adopt some belief just in case either $P(R)$ or

³ Notice that this utility function only reflects *cognitive* value. Certainly there are cases in which due to extra-cognitive values, the ordering is not preserved: suppose that there is some completely innocent symptom that shows up in blood tests, which the medical world typically misdiagnoses and aggressively treats with very hazardous medication. Then the utility of agnosticism clearly exceeds the utility of either true belief or false belief.

⁴ To construct a Von Neumann-Morgenstern utility function, we start off with an ordering over the state of affairs and assign the utility value 1 to the best state of affairs S_1 and the utility value 0 to the worst state of affairs S_n . An intermediate state of affairs S_i receives the utility value a just in case the person in question is indifferent between the state of affairs S_i holding for certain and a lottery with S_1 and S_n as the possible prizes such that $P(S_1) = a$ and $P(S_n) = (1 - a)$. See Luce & Raiffa, (1957: 19-23).

$P(\text{not-}R)$ is greater than or equal to $U(\text{Agnosticism})$ ⁵. Hence, the utility of agnosticism is precisely the threshold for belief acceptance.

The following analogy may be illuminating. One could ask oneself whether taking up a friendship is a good thing when we are not quite certain of the moral character of our would-be friend. The following preference ordering would be quite common: taking up a friendship with a good person ranks above refraining from taking up a friendship, which in turn ranks above taking up a friendship with a bad person. When I ask you about the strength of your preference, you may well wish to differentiate. For instance, for romantic friendships, you may be risk averse, i.e. you may be just about indifferent between taking up a friendship with a good person and refraining from taking up a friendship, while, for political friendships, you may be risk prone, i.e. you may be just about indifferent between refraining from taking up a friendship and taking up a friendship with a bad person. Furthermore different people may profess to different utility functions within a particular area of life.

William James's "The Will to Believe" (1948) is a *locus classicus* of doxastic voluntarism. James argues, against W.K. Clifford (1879), that we have the right to adopt religious beliefs even in the face of insufficient evidence. We could fill in James's argument as follows: religious beliefs are the kind of beliefs for which the cognitive value of agnosticism may only marginally exceed the cognitive value of false belief; in other words, the value of the parameter a may well be in the neighborhood of 0, or at least smaller than or equal to .50. At some junctures, James seems to be saying that we have the freedom to define a risk-prone utility function over the relevant doxastic states, e.g. '...these feelings of our duty about either truth or error are in any case only expression of our passionate life.' (1948, 100) At other junctures, he seems to think that there is some disputing to be done about our tastes for truth and error and that in the case of religious beliefs the weight of the argument is on the side of a risk-prone utility function. Religious belief affords some 'vital good' in life that is lost by agnosticism: the threat of this loss gives us good reason to let the utility value of agnosticism drop into the neighborhood of the utility value for false belief.⁶

⁵ Strictly speaking, we are indifferent between adopting the belief that R and remaining agnostic if and only if $P(R)$ exceeds $P(\text{not-}R)$ and $P(R) = U(\text{Agnosticism})$. We adopt the belief that R if $P(R) > P(\text{not-}R)$ and $P(R) > U(\text{Agnosticism})$ and we remain agnostic if $U(\text{Agnosticism}) > P(R)$ and $U(\text{Agnosticism}) > P(\text{not-}R)$.

⁶ There is a difficult issue lurking here. One might object that if James is concerned about the loss of a vital good, then he is not strictly concerned with cognitive value, and hence James and Clifford are arguing across purposes: While Clifford is talking about epistemic integrity, James is talking about the rationality of belief, which may ride on extra-cognitive values as well (cf. footnote 3). But I think that this objection is mistaken: as long as the goodness that James is alluding to is intrinsic to the believing itself, and is not some consequence of the believing (as in Pascal's Wager), then the value is of a strictly cognitive nature.

The analogy to taking up a friendship is instructive here, and James actually runs a similar comparative argument about a man who 'hesitate[s] indefinitely to ask a certain woman to marry him because he was not perfectly sure that she would prove an angel after he brought her home.' (1948, 106) One might certainly argue for the freedom to hold a more risk-prone utility function when it comes to taking on romantic friendships and be more light-hearted about falling in love, even if one is not entirely sure about the moral character of the would-be lover. And one might even argue for such a life style by saying that it would be a mistake to forego the possibility of a 'vital good' like love in one's life. But of course the argument would only carry so far: just as Clifford may reaffirm his personal abhorrence for error, a person who is more cautious about love may reaffirm his personal abhorrence for being stuck in a relationship with a person of a dubious moral character.

On the Lockean notion of belief, qualitative doxastic states supervene on a base which does not only include quantitative doxastic states but also particular information about the person's cognitive values, viz. how strongly does he prefer having true beliefs over being agnostic and being agnostic over having false beliefs within a certain area of inquiry. Quantitative doxastic states in conjunction with these cognitive values determine the qualitative doxastic state of the agent. Note that the failure of the supervenience of qualitative on quantitative doxastic states only permits a weak difference on this notion of belief. It is possible for two persons to be in the same quantitative doxastic state, and yet one person believes that R while the other does not believe that R. What is not possible is that one person believes that R, while the other person believes that not-R. I will show that supervenience not only fails for the remaining notions of belief, but also that this failure will leave room for a strong difference in qualitative doxastic states.

4 THE AGENTIC NOTION OF BELIEF

On the agentic notion of belief it is easy to see why the supervenience thesis fails. Suppose you and I both assign the same relatively low degree of confidence c to there being a witch in the woods. We both prefer Sunday walks in the woods without a witch to Sunday walks in suburbia to Sunday walks in the woods with a witch. Hence, our utility functions will have the following structure:

$$\begin{aligned} U(\text{Walk-in-Woods} \mid \text{No-Witch-in-Woods}) &= 1 \\ U(\text{Walk-in-Suburbia}) &= a^7 \end{aligned}$$

⁷ One might object that the opportunity cost of a walk in suburbia while there's no witch in the woods is greater than the opportunity cost of a walk in suburbia while there's a witch in the woods. Fair enough, but this complication changes nothing substantially as long as we make the plausible assumption that the following preference ordering holds: walk in the woods without a witch > walk in suburbia while there is a witch in the

(continued)

$$U(\text{Walk-in-Woods} \mid \text{Witch-in-Woods}) = 0$$

Now suppose that, while your concern that there is a witch in the woods does not stop you from taking Sunday walks in the woods, it does stop me: You are much less worried about an encounter with a witch than I am, and furthermore, you derive much greater enjoyment from your Sunday walks in the woods rather than in suburbia than I do. Then whereas the value of the parameter a will be close to 1 for me, it is close to 0 for you. So calculating the expected utility of the available courses of action:

$$EU(\text{walking-in-the-woods}) = P(\text{Witch-in-Woods})U(\text{Walk-in-Woods} \mid \text{Witch-in-Woods}) + P(\text{No-Witch-in-Woods})U(\text{Walk-in-Woods} \mid \text{No-Witch-in-Woods}) = c$$

$$EU(\text{walking-in-suburbia}) = U(\text{Walk-in-Suburbia}) = a$$

it becomes clear why you are to be found in the woods and I am to be found in suburbia on our Sunday walks: although we share the same value for the parameter c , we do not share the same value for the parameter a :

$a > c$ for me, while $c > a$ for you.

Philosophers may protest at this point that this is at best a very quirky notion of belief. If you and I think it more likely than not that there is no witch in a woods, then one may at best say that I believe that there *might* be a witch in the woods, and since we both assign the same degree of confidence to there being a witch in the woods, consistency demands that one also say that you believe that there *might* be a witch in the woods. Fair enough: there is indeed some notion of belief on which this is precisely what we should say. But at the same time, the agentic notion of belief is not an uncommon notion in everyday life. Consider a person who leads his life in such a way that many of his actions are properly explained by his desire to avoid an encounter with a witch.⁸ His degree of confidence that witches exist may even be quite low, but he is so engaged with the matter that we commonly do ascribe the belief to him that witches exist. In fact, his degree of confidence may not be any higher than the degree of confidence of someone whose actions never require any explanation by her desire to avoid an encounter with a witch. And it would be equally common to say about such a person that although she does not rule out the existence of witches, she does not believe that

woods > walk in suburbia while there is no witch in the woods > walk in the woods with a witch.

⁸ Notice that this is not inconsistent with saying that the person's degree of confidence that there is a witch in the woods is low, since explanation targets his *actual* actions, while ascriptions of degrees of confidence rest on his disposition to act (e.g. his betting behavior) under certain *counterfactual* conditions. For an alternative procedure to ascribe degrees of confidence that does not rest on betting behavior and hence is independent of the person's utility function, see Hawthorne and Bovens (1999, 254-255 and 260-261).

witches exist. I see no reason why a philosophical account should correct such a practice of belief ascription, rather than leave room for a notion of belief which captures this practice.

Another *locus classicus* of doxastic voluntarism is Pascal's Wager. The agentic notion of belief permits us to take some of the paradoxical features out of the Wager. We construct the following plausible utility function:

$$\begin{aligned} U(\text{Believing} \mid \text{God-Exists}) &= U(\text{Eternal-Bliss}) = 1 \\ U(\text{Not-Believing} \mid \text{God-Does-Not-Exist}) &= a + \varepsilon, \text{ for } 0 < a + \varepsilon < 1 \\ U(\text{Believing} \mid \text{God-Does-Not-Exist}) &= a, \text{ for } 0 < a < 1 \\ U(\text{Not-Believing} \mid \text{God-Exists}) &= U(\text{Eternal-Damnation}) = 0 \end{aligned}$$

Eternal bliss is the best outcome, while eternal damnation is the worst outcome. In between are outcomes in which there is no after-life to be reckoned with. Now if God does not exist, then it is presumably still better to have had true beliefs rather than to have had false beliefs. What we are reminded of, is that relative to the outcomes of eternal bliss and eternal damnation⁹, this difference, which is a difference in an earthly good, must become very miniscule. Setting my subjective probability for the existence of God at c , we can calculate the rational course of action:

$$\begin{aligned} EU(\text{believing}) &= U(\text{Believing} \mid \text{God-Exists})P(\text{God-Exists}) + \\ &\quad U(\text{Believing} \mid \text{God-Does-Not-Exist})P(\text{God-Does-Not-Exist}) \\ &= c + a(1 - c) = c + a - ac \end{aligned}$$

$$\begin{aligned} EU(\text{not-believing}) &= U(\text{not-Believing} \mid \text{God-Exists})P(\text{God-Exists}) + \\ &\quad U(\text{not-Believing} \mid \text{God-Does-Not-Exist})P(\text{God-Does-Not-Exist}) \\ &= (a + \varepsilon)(1 - c) = a - ac + \varepsilon - \varepsilon c \end{aligned}$$

The expected utility of *believing* exceeds the expected utility of *not-believing* if and only if $c > \varepsilon - \varepsilon c$, and consequently if $c > \varepsilon$. So unless we have certainty that God does not exist (and what could ground such certainty?), it is rational to believe, since 0 is very miniscule. What is paradoxical about the wager is this: how can we believe when our degree of confidence in the existence of God is low? On the agentic notion of belief, the paradox vanishes. By not going into the woods on Sunday, we may be said to believe that there is a witch in the woods, even if we think it more likely than not that there is no witch in the woods. Similarly, by doing all the things that a believer does, we may be said to believe that God exists, and we may self-ascribe such a belief, even if we continue to think it more likely than not that God does not exist. Hence, we can remain Lockean non-

⁹ Pascal actually capitalizes only on eternal bliss rather than on both eternal bliss and eternal damnation to make the point in question.

believers and at the same time, on Pascal's counsel, become agentic believers.¹⁰ Of course, whether God (assuming that he exists and that he is the kind of God that rewards religious belief) finds mere agentic belief sufficient to award eternal bliss remains unknown to us!¹¹

On the agentic notion of belief, qualitative doxastic states supervene on a base which does not only include quantitative doxastic states but also particular information about the person's extra-cognitive values, viz. his preferences for states of the world that would ensue from his actions, conditional on the value of the propositional variable in question. Does sameness of quantitative doxastic states permit a weak or a strong difference between qualitative doxastic states? Let us return to the response: "I believe enough not to go up there." Suppose that our speaker had said something weaker, viz. "I believe enough not to go up there on Halloween night, but otherwise, it doesn't bother me." Or suppose he had said something stronger, viz. "I believe enough to sell my house promptly and to move out of the state." One's place on the continuum does not need to be determined by one's degree of confidence that there is a witch in the woods: we can fix the degree of confidence at some non-extreme value, and explain these attitudes fully in reference to some difference in one's fear of witches. What is witty about the expression "I believe enough (not) to..." is that it mixes qualitative and quantitative doxastic talk. When we are forced to use straight qualitative talk, we would expect the following pattern of belief attribution on the agentic notion: If one's actions are minimally governed by a concern for witch-avoidance, one is said to believe that there's no witch in the woods; If they are somewhat governed by a concern for witch-avoidance, one is said neither to believe that there's a witch in the woods, nor to disbelieve it; If they are maximally governed by a concern for witch-avoidance, one is said to believe that there's a witch in the woods. Since all these belief attributions are consistent with a fixed degree of confidence, the failure of supervenience permits both a weak and a strong difference in qualitative doxastic states.

5 THE ABDUCTIVE NOTION OF BELIEF

Let us turn straightaway to an example. Suppose that a patient comes in to a hospital for a routine check. The positive values of the propositional variables in S

¹⁰ I am not claiming to be doing textual exegesis here. Pascal enjoins us to start acting as if we are already believing and he trusts that our actual beliefs will follow suit. Certainly, it is possible that faced with self-induced cognitive dissonance, we could become Lockean believers as well. But although we can admit to the psychological possibility, the objection remains that this course of action violates normative constraints on belief. My suggestion is that the objection can be avoided if the Pascalian route to belief is conceived of as a route to agentic belief, not to Lockean belief.

¹¹ Note however my pessimism in footnote 16.

(i) ascribe certain diseases to this patient, (ii) ascribe certain actions to the patient that may have caused or prevented these diseases, or (iii) ascribe certain symptoms to the patient. There are intricate probabilistic dependencies between these propositional variables which are embedded in the prior probability distribution over these variables. As we gain more information about the patient through questioning, laboratory tests,... we can update the probability distribution.¹² Now suppose that I have gathered evidence about all but n variables in S and that the posterior probability distribution over these n variables has the following form:

$P(R_1, R_2, \dots, R_n) = a$ for a being close to but smaller than .50

$P(R_1, R_2, \dots, R_n) = (1-a) / (2^n - 1)$ for all other $2^n - 1$ combinations of values of the variables R_1, R_2, \dots, R_n

For a sufficiently high parameter n ¹³, it would not be uncommon to say that we believe that R_1, R_2, \dots, R_n , since this belief set contains the most probable combination of propositions. Just like the paleontologist we are inclined to believe the most probable story about some relevant part of the world, even if this story is more likely to be false than true, on the ground that any alternative story is much less probable.

To make the abductive notion of belief¹⁴ palatable, we have insisted that the most probable combination of values is much more probable than any one of the other combinations, and furthermore that, of these less probable combinations, there is no basis to favor one over the other. But in laying out a general definition, I am a bit less colorful than in my advertisement: a person believes abductively that R , just in case R is contained in the most plausible combination of values of the relevant set of propositional variables. We will return to the question as to how we can determine the relevant set of propositional variables.¹⁵

¹² Such scenarios are standardly modelled in the theory of Bayesian Networks in artificial intelligence. The example is inspired by an example in Lauritzen and Spiegelhalter (1988), which is implemented in a demo-network of the HUGIN group, entitled 'Asia'. See Jensen, F. (1996) and also <http://www.hugin.com>.

¹³ For example, suppose that I have not received evidence about five propositional variables, R_1, R_2, \dots, R_5 . The posterior joint probability $P(R_1, R_2, \dots, R_5)$ is lower than but in the neighborhood of .50, say .49, and the posterior joint probability of each of the other 31 possible combination of values is $P(\text{not-}R_1, R_2, \dots, R_5) = P(R_1, \text{not-}R_2, \dots, R_5) = \dots = P(\text{not-}R_1, \text{not-}R_2, \dots, \text{not-}R_5) = (1-a) / (2^5 - 1) = .51/31 = .0165$.

¹⁴ I am using C.S. Peirce's notion of *abduction* more broadly than it was intended, since I do not require that the variables, of which I determine what combination of values is most likely given the evidence, are *explanatory* variables. For instance, one could determine the most likely combination of the values of the symptom variables, given that one has acquired evidence about what diseases are afflicting the patient.

¹⁵ I leave it as an open question whether one could make good in some principled way on the color that was promised in my advertisement and what sort of claims one could make about the abductive notion of belief under such more severe constraints.

Notice that the abductive notion of belief may conflict with the Lockean notion of belief. Suppose that we have gathered evidence about all but 2 variables in S and that the posterior probability distribution over these variables is as follows:

$$\begin{aligned} P(R_1, R_2) &= .20 \\ P(\text{not-}R_1, R_2) &= .39 \\ P(R_1, \text{not-}R_2) &= .41 \\ P(\text{not-}R_1, \text{not-}R_2) &= 0 \end{aligned}$$

so that the marginal distributions are $P(R_1) = .61$ and $P(R_2) = .59$. Hence for a sufficiently low threshold, the Lockean notion of belief enjoins us to adopt the belief set $\{R_1, R_2\}$, whereas the abductive notion of belief enjoins us to adopt $\{R_1, \text{not-}R_2\}$.

What makes a set of propositional variables relevant? Whether we form beliefs about some propositional variable or other depends on our cognitive interest. Suppose that I know that 90% of Americans consume at least one beer per week. As I open up the phone book my eye falls on the name Paul Smith. My degree of confidence that Paul Smith has had at least one beer in the past week is .90, but do I believe that Paul Smith has had a beer? I don't think so. The problem is not that I have a higher threshold for the matter at hand. Rather, I simply have too little of a cognitive interest in whether Paul Smith has had a beer or not to form any beliefs about Paul Smith's beer drinking habits.

Let us return to the hospital. Suppose that a medical doctor, a nurse and a detective all share the same probability distribution over the propositional variables that are pertinent to our patient's health, i.e. over all the variables in S . As new evidence comes in, they update their probability distributions in the same fashion. What interests the medical doctor is what disease or combination of diseases the patient suffers from. What interests the nurse is what combination of symptoms he will have to deal with during his night shift. What interests the detective is what combination of causes can provide information about the patient's habits and past actions. On the abductive notion of belief, the medical doctor will believe the most probable combination of values of the disease variables, the nurse will believe the most probable combination of values of the symptom variables, and the detective will believe the most probable combination of values of the cause variables. They will pay attention to how variables that are outside their range of cognitive interest affect the variables inside their range of cognitive interest, but they are no more prone to form beliefs about variables outside their own range than I am to form a belief about whether Paul Smith has had a beer last week.

The data are on the table to make the supervenience thesis crash for the abductive notion of belief. In the hospital case, we have partitioned the relevant variables according to the respective interests of the medical doctor, the nurse and the detective. But of course, nothing stands in the way of two people sharing some, but

not other cognitive interests. So let us assume that two people share the same probability distribution over a set of propositional variables, and that one person takes a cognitive interest in the variables R_1 and R_2 , while the other person takes a cognitive interest in R_2 and R_3 . Furthermore, they share the following probability distribution over $R_1 R_2 R_3$:

$$P(R_1, R_2, R_3) = .20$$

$$P(R_1, \text{not-}R_2, R_3) = .20$$

$$P(R_1, \text{not-}R_2, \text{not-}R_3) = .30$$

$$P(\text{not-}R_1, R_2, R_3) = .30$$

$$P(R_1, R_2, R_3) = 0 \text{ for any other combination of values of the variables } R_1, R_2, R_3.$$

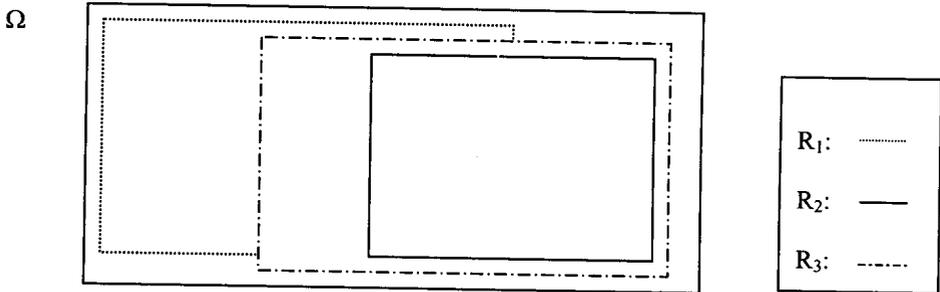


Figure 1: The Probability Space over the Variables R_1 , R_2 and R_3

This probability distribution is also represented in figure 1. We can calculate the following marginal distributions for $R_1 R_2$ and $R_2 R_3$:

$$P(R_1, R_2) = .20$$

$$P(R_1, \text{not-}R_2) = .50$$

$$P(\text{not-}R_1, R_2) = .30$$

$$P(\text{not-}R_1, \text{not-}R_2) = 0$$

$$P(R_2, R_3) = .50$$

$$P(R_2, \text{not-}R_3) = 0$$

$$P(\text{not-}R_2, R_3) = .20$$

$$P(\text{not-}R_2, \text{not-}R_3) = .30$$

On the abductive notion of belief, a person with a cognitive interest in R_1 and R_2 will form the belief set $\{R_1, \text{not-}R_2\}$, while a person with a cognitive interest in R_2 and R_3 will form the belief set $\{R_2, R_3\}$. Although they share precisely the same quantitative doxastic states, one will come to believe that R_2 while the other one will come to believe that $\text{not-}R_2$. The lesson to be learned here is that on the abductive notion of belief, qualitative doxastic states do not supervene on quantitative doxastic states, but rather on quantitative doxastic states conjoint with cognitive interests. In the hospital case, the failure of supervenience permits weak differences, i.e. two persons may be in the same quantitative doxastic states with respect to the variables R_1, R_2, \dots, R_n , but, for some $i=1, \dots, n$, one person comes to believe R_i , while the other person does not come to believe R_i , since R_i is not within her cognitive interests. The latter case shows that the failure of supervenience permits strong differences, i.e. two persons may be in the same quantitative

doxastic states with respect to the variables R_1, R_2, \dots, R_n , but one person comes to believe R_i while the other person comes to believe not- R_i for some $i=1, \dots, n$.

6 REPERCUSSIONS FOR VOLUNTARISM AND EVIDENTIALISM

There is good news and there is bad news in my analysis for the doxastic voluntarist, though I fear that the bad news overshadows the good news. Let's start with the good news. My analysis shows that there are notions of belief in which valuational attitudes do play an important role. Different kinds of valuational attitudes play a role: the Lockean notion rides on cognitive value, the agentic notion rides on extra-cognitive value and the abductive notion rides on cognitive interests. Some voluntarists might welcome this. Voluntarists of a more existentialist flavor will insist on something stronger than this, viz. that our beliefs are under the direct control of the will. Now, in so far that the relevant values are under the control of the will, my analysis may have something to offer to them as well. But whether our values are indeed under the control of the will is, of course, a separate discussion. I have argued elsewhere (1995) that it may be reasonable to intentionally change one's preferences, either directly or in some roundabout manner, but not one's moral values. This imposes limitations on how much my analysis has to offer to the existentialist voluntarist, since some of the relevant valuational attitudes are not of a preferential nature.

But now for the bad news. Somehow I have this nagging feeling that voluntarist are really not finding much joy in what I have to offer. So why is that, after turning my back on the evidentialists, I am not making any friends in the voluntarist camp? Although this is by no means uncontroversial, supervenience theses are often thought to have ontological repercussions: the supervening properties are thought to be ontologically suspect, while only the base properties survive Ockham's razor. In Armstrong words, what supervenes is 'an ontological free lunch', the 'truth-makers, the ontological ground' are to be found amongst the base-properties. (Armstrong, 1989: 56) In this vein, one might say that our cognitive states *proper* are our quantitative doxastic states: belief talk is just window-dressing of various brands, that stands in complex relationships to our degrees of confidence conjoint with all kinds of valuational attitudes. The voluntarist has enjoyed a short moment of glory: there are indeed many notions of belief in which valuational attitudes play an important role. But the price was certainly high: it is degree-of-confidence talk and not belief talk that ought to be spoken on the doxastic stage *proper*.¹⁶ My suspicion is that if I really want to curry favor amongst voluntarists, I would have to admit that doxastic rationality permits that degrees of confidence themselves are legitimately affected by valua-

¹⁶ Notice that if supervenience really makes beliefs less than fully respectable entities, then there is no reason to think that God would care much about our agentic beliefs. Hence Pascal's Wager would be of little consequence, as long as we do not succeed in increasing our degree of confidence that God exists.

tional attitudes. I am afraid that I cannot oblige: on the doxastic stage *proper*, I find myself entirely comfortable in the evidentialist camp.¹⁷

REFERENCES

- Armstrong, D. (1989) *Universals – an Opinionated Introduction*, (Boulder: Westview Press, CO).
- Bovens, L. (1995) “The Intentional Acquisition of Mental States”, *Philosophy and Phenomenological Research*, 40, 821-840.
- Clifford, W.K. (1879) *Lectures and Essays* (London: Macmillan), excerpted in Pojman, L. P. (1998) *Philosophy of Religion: an Anthology*. (Wadsworth, Belmont, CA), 399-403.
- Foley, R. (1992) “The Epistemology of Belief and the Epistemology of Degrees of Belief”, *American Philosophical Quarterly* 29, 111-121.
- Hawthorne, J. & Bovens, L. (1999) “The Preface, the Lottery and the Logic of Belief”, *Mind* 108, 241-264.
- James, W. (1948) “The Will to Believe”, in *Essays on Pragmatism*, ed. by A. Castell (New York: Hafner), 88-109. Originally published in the *New World*, June 1869.
- Jensen, F.V. (1996) *An Introduction to Bayesian Networks*. (Springer, Berlin).
- Lauritzen, S.L. and Spiegelhalter, D.J. (1988) “Local Computations with Probabilities on Graphical Structures and their Application to Expert Systems” (with Discussion), *Journal of the Royal Statistical Society, Series B* 50, 157-224.
- Luce, R.D. and Raiffa, H. (1957) *Games and Decisions*. (New York: Wiley & Sons)
- Pascal, B. (1910) *Thoughts*, translated by W.F. Trotter (New York: Collier & Son), excerpted in Pojman, L. P. (1998) *Philosophy of Religion: an Anthology*. (Wadsworth, Belmont, CA), 397-399.
- Sanchez, E. and Myrick, D., directors (1999) *The Blair Witch Project*.

¹⁷ I thank Michael Huemer, Iain Martel, Wes Morriston, Graham Oddie, Michael Tooley and the participants of the conference on ‘Belief and the Will’ in Nijmegen, March 1999, for comments and discussion. My research was in part conducted at the University of Konstanz, supported by the Alexander von Humboldt Foundation.