

RATIONALITY AND SELF-CONFIDENCE

Frank Arntzenius, Rutgers University

1. Why be self-confident?

Hair-Brane theory is the latest craze in elementary particle physics. I think it unlikely that Hair-Brane theory is true. Unfortunately, I will never know whether Hair-Brane theory is true, for Hair-Brane theory makes no empirical predictions, except regarding some esoteric feature of the microscopic conditions just after the Big Bang. Hair-Brane theory, obviously, has no practical use whatsoever. Still, I care about the truth: I want my degree of belief $D(H)$ in Hair-Brane theory H to be as close to the truth as possible. To be precise:

- (1) if H is true then having degree of belief $D(H)=r$ has epistemic utility $U(D)=r$ for me
- (2) if H is false then having degree of belief $D(H)=r$ has epistemic utility $U(D)=1-r$ for me.

Currently, my degree of belief $D(H)=0.2$. Am I, by my own lights, doing a good epistemic job? Let's see. The expected epistemic utility EU of degree of belief $D'(H)=r$, given my current degree of belief D , is:

$$(3) EU(D')=D(H)U(D' \& H)+D(\neg H)U(D' \& \neg H)=0.2r+0.8(1-r)=0.8-(0.6)r.$$

Obviously, $EU(D')$ is maximal for $D'(H)=0$. So, by my own lights, I am not doing a good job; I would do better if I were absolutely certain that Hair-Brane theory is false. Unfortunately, I am not capable of setting my degrees of belief at will. All I can do is recognize my own epistemic shortcomings. So I do.

The above is a strange story. Real people, typically, do not judge their own degrees of belief as epistemically deficient. To coin a term: real people tend to be 'self-confident'. The puzzle that Gibbard poses is that he can see no good reason to be self-confident. For, according to Gibbard, all that follows from having the truth as one's goal, all that follows from having the accuracy of one's state of belief as one's desire, is that a higher degree of belief in the truth is better than a lower degree of belief in the truth. That is to say, according to Gibbard, the only constraint on the epistemic utilities of a rational person is that they should increase as her degrees of belief get closer to the truth. The simplest, most natural, epistemic utility function ('scoring' function) which satisfies this constraint, is a linear function. In the case of a single proposition, the function that I stated in (1) and (2) is such a function. So, according to Gibbard, not only is it rationally acceptable to judge one's own degrees of belief as epistemically deficient, it is very natural to do so.

In the next section I will suggest that considerations regarding updating can serve to explain why real people are self-confident. However, I will then go on to explain why I am nonetheless sympathetic to Gibbard's suggestion that one can not give a purely epistemic justification for why our belief states are as they are.

2. Updating and self-confidence

Gibbard's considerations are entirely synchronic. That is to say, he does not consider the evolution of one's belief state through time. But having the truth as one's goal surely includes the desire to get closer to the truth as time passes. In this section I will try to incorporate such

considerations.

Let's start with a simple example. Suppose I initially have the following degree of belief distribution D:

- (4) $D(H\&E)=0.4$
- (5) $D(H\&\neg E)=0.2$
- (6) $D(\neg H\&E)=0.1$
- (7) $D(\neg H\&\neg E)=0.3$

And suppose that I have a linear epistemic utility function. In particular, suppose that, according to my current degrees of belief D, the expected epistemic utility of degree of belief distribution D' is:

$$(8) 0.4D'(H\&E)+0.2D'(H\&\neg E)+0.1D'(\neg H\&E)+0.3D'(\neg H\&\neg E)$$

This is maximal for $D'(H\&E)=1$. So, epistemically speaking, I desire that I currently be certain that H&E is true, even though in fact I am not certain of that at all.

Now suppose that I know that in one hour I will learn whether E is true or not. And suppose that the only thing I now care about is the degrees of belief that I will have one hour from now. If that is so, what should I now regard as epistemically the best policy for updating my degrees of belief in the light of the evidence that I will get? That is to say, what degrees of belief D_E do I now think I should adopt if I were to get evidence E, and what degrees of belief $D_{\neg E}$ do I now think I should adopt if I were to get evidence $\neg E$? Well, my current expected epistemic utility for my future degrees of belief is:

$$(9) 0.4U(H\&E\&D_E)+0.2U(H\&\neg E\&D_{\neg E})+0.1U(\neg H\&E\&D_E)+0.3U(\neg H\&\neg E\&D_{\neg E}).$$

We can expand each of the four epistemic utilities that occur in (9):

- (10) $U(H\&E\&D_E)=D_E(H\&E)+(1-D_E(H\&\neg E))+(1-D_E(\neg H\&E))+(1-D_E(\neg H\&\neg E))$
- (11) $U(H\&\neg E\&D_{\neg E})=(1-D_{\neg E}(H\&E))+D_{\neg E}(H\&\neg E)+(1-D_{\neg E}(\neg H\&E))+(1-D_{\neg E}(\neg H\&\neg E))$
- (12) $U(\neg H\&E\&D_E)=(1-D_E(H\&E))+(1-D_E(H\&\neg E))+D_E(\neg H\&E)+(1-D_E(\neg H\&\neg E))$
- (13) $U(\neg H\&\neg E\&D_{\neg E})=(1-D_{\neg E}(H\&E))+(1-D_{\neg E}(H\&\neg E))+(1-D_{\neg E}(\neg H\&E))+D_{\neg E}(\neg H\&\neg E)$

After substituting these terms into (9) and fiddling around a bit we find that my expected epistemic utility is:

$$(14) 3+0.3D_E(H\&E)-0.5D_E(H\&\neg E)-0.3D_E(\neg H\&E)-0.5D_E(\neg H\&\neg E)-0.5D_{\neg E}(H\&E)-0.1D_{\neg E}(H\&\neg E)-0.5D_{\neg E}(\neg H\&E)+0.1D_{\neg E}(\neg H\&\neg E).$$

This expression is maximized by setting $D_E(H\&E)=1$ and $D_{\neg E}(\neg H\&\neg E)=1$ (and setting the other degrees of belief equal to 0). So if all I care about is the degrees of belief I will have one hour from now, then I should update on E by becoming certain that H&E is true, and I should update on $\neg E$ by becoming certain that $\neg H\&\neg E$ is true. In particular, by my current lights, it would be wrong to first change my degrees of belief so as to maximize my current expected epistemic utility, then update these degrees of belief by conditionalization, and then change these

conditionalized degrees of belief so as to maximize expected epistemic utility by the lights of these conditionalized degrees of belief. So there is a conflict between maximizing the expected epistemic utility of my current degrees of belief (by my current lights), and maximizing the expected epistemic utility of my future degrees of belief (by my current lights). At least there is such a conflict, if I update by conditionalization.

Given that there is such a purely epistemic conflict, the obvious question is: what should I do if I only have epistemic concerns and I care both about the accuracy of my current degrees of belief and about the accuracy of my future degrees of belief? One might answer: no problem, I should maximize expected epistemic utility (by my current lights) of both of my current degrees of belief and my future degrees of belief, and hence I should jettison conditionalization. That is to say I should now set my degree of belief to $D'(H \& E) = 1$. And then, if I get evidence E , my degrees of belief should stay the same, but if I get evidence $\neg E$, I should set my degrees of belief to $D_{\neg E}(\neg H \& \neg E) = 1$. Unfortunately, there are two problems with this answer.

In the first place, it seems worrying to jettison conditionalization. The worry is not just the general worry that conditionalization is part of the standard Bayesian view. The worry, more specifically, is that if one rejects conditionalization one will have to reject standard arguments in favor of conditionalization, namely diachronic Dutch books arguments. But if one does that, shouldn't one also reject synchronic Dutch book arguments? And if one does that, then why have degrees of belief, which satisfy the axioms of probability, to begin with? I will return to this question in section 4. For now, let me turn to the second problem.

The second problem is that if one were to re-set one's current degrees of belief so as to maximize one's current expected epistemic utility, one would thereby lose the ability to set one's future degrees of belief so as to maximize the current expected epistemic utility of those future degrees of belief. Let me explain this in a bit more detail.

According to my current degrees of belief D the epistemically best current degree of belief distribution is:

$$(15) D'(H \& E) = 1$$

$$(16) D'(H \& \neg E) = 0$$

$$(17) D'(\neg H \& E) = 0$$

$$(18) D'(\neg H \& \neg E) = 0$$

Now, according to my original plan, if I were to learn E then I should update by becoming certain that $H \& E$ is true, and if I were to learn $\neg E$ then I should become certain that $\neg H \& \neg E$ is true. But if I were to replace D by D' then I would lose the information as to what I should do were I to learn $\neg E$. The reason why I originally desire to update on $\neg E$ by becoming certain that $\neg H \& \neg E$, rather than becoming certain that $H \& \neg E$, is that $D(\neg H / \neg E)$ is higher than $D(H / \neg E)$. But if I were to change D into D' the relevant information is no longer encoded in my degrees of belief: D' could have come from a degree of belief D (via expected epistemic utility maximization) according to which $D(\neg H / \neg E)$ is lower than $D(H / \neg E)$, but it could also have come from one according to which $D(\neg H / \neg E)$ is higher than $D(H / \neg E)$. That is to say, if one's epistemic utilities are linear, then maximizing the expected epistemic utility (by one's current lights) of one's degrees of belief can make it impossible to maximize the expected epistemic utility (by one's current lights) of one's degrees of belief at a future time.

The obvious solution to this problem is for the ideal rational agent to have two separate degree of belief distributions. An ideal rational agent should have a 'prudential' degree of belief

distribution, which she uses to guide her actions and to compute epistemic utilities, and an ‘epistemic’ degree of belief distribution, which she always sets in order to maximize epistemic utility.

Now, one might worry that there is still going to be a problem. For consider again the example that I started this section with, i.e. suppose that my initial prudential degrees of belief are

- (19) $D^{pr}(H\&E)=0.4$
- (20) $D^{pr}(H\&\neg E)=0.2$
- (21) $D^{pr}(\neg H\&E)=0.1$
- (22) $D^{pr}(\neg H\&\neg E)=0.3$

Suppose I use these initial prudential degrees of belief to set my initial epistemic degrees of belief so as to maximize expected epistemic utility. Then my initial epistemic degrees of belief would be:

- (23) $D^{ep}(H\&E)=1$
- (24) $D^{ep}(H\&\neg E)=0$
- (25) $D^{ep}(\neg H\&E)=0$
- (26) $D^{ep}(\neg H\&\neg E)=0$

Now I don’t (yet) need to worry that I have lost the possibility of maximizing the expected epistemic utility (according to my initial prudential degrees of belief) of my epistemic degrees of belief one hour from now, since, even though I adopted initial epistemic degrees of belief as indicated, I have retained my initial prudential degrees of belief. However there might still be a problem. For when I acquire evidence E, or evidence $\neg E$, I will, presumably, update my prudential degrees of belief by conditionalization. So will our problem therefore re-appear? Will my updated prudential degrees of belief contain enough information for me to be able to deduce from them which epistemic degree of belief distribution has maximal expected epistemic utility according to my initial prudential degree of belief distribution? And, even if I do have enough information to be able to stick to my original plan, will that plan still look like a good plan according to my updated prudential degrees of belief? Let’s see.

Recall that according to my initial prudential degrees of belief, if all I care about is the epistemic utility of my degrees of belief one hour from now, then I should update on E by becoming certain that H&E is true, and I should update on $\neg E$ by becoming certain that $\neg H\&\neg E$ is true. Now, if I were to learn E and update my prudential degree of belief by conditionalization then my prudential degrees of belief would become

- (27) $D^{pr}(H\&E)=0.66$
- (28) $D^{pr}(H\&\neg E)=0.33$
- (29) $D^{pr}(\neg H\&E)=0$
- (30) $D^{pr}(\neg H\&\neg E)=0$

According to these prudential degrees of belief expected epistemic utility is maximized by being certain that H&E is true.

Similarly, if I were to learn $\neg E$ and I conditionalized on this, then my prudential degrees of belief would become

- (31) $D^{pr}(H \& E) = 0$
(32) $D^{pr}(H \& \neg E) = 0$
(33) $D^{pr}(\neg H \& E) = 0.75$
(34) $D^{pr}(\neg H \& \neg E) = 0.25$

According to these prudential degrees of belief expected epistemic utility is maximized by being certain that $\neg H \& \neg E$ is true.

So, in this case at least, the epistemic degrees of belief that I should adopt in the light of evidence, according to my initial prudential degrees of belief, are the same as the ones that I should adopt according to my later prudential degrees of belief, if I update my prudential degrees of belief by conditionalization.

What it is more interesting, and perhaps more surprising, is that this is true for every possible initial prudential degree of belief distribution, and for every possible epistemic utility function. That is to say: no matter what one's epistemic utilities are, if according to one's prudential degrees of belief at some time t , plan P for updating one's epistemic degrees of belief maximizes expected epistemic utility, then, after one has updated one's prudential degrees of belief by conditionalization, plan P will still maximize expected utility according to one's updated prudential degrees of belief. The proof of this fact for the general finite case is simple, so let me give it.

Let $D^{pr}(W_i)$ be my initial prudential degree of belief distribution over possibilities W_i .¹ Let $U(W_i \& D^{ep})$ be my epistemic utility for having degree of belief distribution D^{ep} in possibility W_i . Suppose I know that in an hour I will learn which of E_1, E_2, \dots, E_n is true (where the E_i are mutually exclusive and jointly exhaustive) An 'epistemic plan' P is a map from current prudential degree of belief distributions plus evidence sequences to future epistemic degree of belief distributions. Let P map D^{pr} plus E_i to D^{ep}_i . Then P has maximal expected epistemic utility according to D^{pr} and U iff for every alternative plan P' (which maps D^{pr} plus E_i to D^{ep}_i') we have:

$$(36) \sum_i \sum_k D^{pr}(W_k \& E_i) U(W_k \& E_i \& D^{ep}_i) \geq \sum_i \sum_k D^{pr}(W_k \& E_i) U(W_k \& E_i \& D^{ep}_i')$$

We can re-write this as

$$(37) \sum_i \sum_k D^{pr}(E_i) D^{pr}(W_k/E_i) U(W_k \& E_i \& D^{ep}_i) \geq \sum_i \sum_k D^{pr}(E_i) D^{pr}(W_k/E_i) U(W_k \& E_i \& D^{ep}_i').$$

The left hand side being maximal implies that each separate i-term is maximal:

$$(38) \sum_k D^{pr}(E_i) D^{pr}(W_k/E_i) U(W_k \& E_i \& D^{ep}_i) \geq \sum_k D^{pr}(E_i) D^{pr}(W_k/E_i) U(W_k \& E_i \& D^{ep}_i'), \text{ for each } i.$$

Therefore

$$(39) \sum_k D^{pr}(W_k/E_i) U(W_k \& E_i \& D^{ep}_i) \geq \sum_k D^{pr}(W_k/E_i) U(W_k \& E_i \& D^{ep}_i'), \text{ for each } i.$$

¹ I am assuming that my degrees of belief are not part of the possibilities W_i that I distribute my degrees of belief over.

But this just means that if we conditionalize D^{pr} on E_i then, according to the resulting degree of belief distribution (and U), the expected epistemic utility of D^{ep}_i is maximal.

Let me now summarize what we have seen in this section, and draw a tentative conclusion.

No matter what one's epistemic utility function is, one can maximize one's epistemic utilities at all times by having two separate degree of belief distributions: a prudential degree of belief distribution which guides one's actions and one's choice of an epistemic degree of belief distribution, and an epistemic degree of belief distribution whose sole purpose is to maximize epistemic utility. One can then give a purely epistemic argument for updating one's prudential degrees of belief by conditionalization, on the grounds that such updating guarantees cross-time consistency of epistemic utility maximization. The epistemic degrees of belief of an ideal agent at a given time do not determine how she updates her epistemic degrees of belief in the light of evidence. Rather, she updates her epistemic degrees of belief by first conditionalizing her prudential degrees of belief and then maximizing epistemic utility. Thus the epistemic degrees of belief of an ideal agent are largely epiphenomenal: they are only there to maximize the 'epistemic score' of an agent, they are not there to guide her actions, nor are they there to help determine her future epistemic degrees of belief. This suggests that rational people can make do without epistemic utilities and epistemic degrees of belief, which could explain why real people do not consider themselves epistemically deficient. Let me bolster this suggestion by arguing that it is not clear what epistemic utilities are.

3. What are epistemic utilities?

Gibbard characterizes epistemic utilities, roughly, as follows. Person P 's epistemic utilities are the utilities that P would have were P to ignore both the 'guidance' value and the 'side' values of his degrees of belief. The 'guidance' value of P 's degrees of belief is the value these degrees of belief have for P due to the way in which they guide P 's actions. The 'side' values of P 's degrees of belief for P are values such as P 's happiness due to, e.g., P 's certitude that he will have a pleasant afterlife, or P 's dejection due to e.g. P 's certitude of his own moral inferiority, and so on. My worry now is that it is not clear what epistemic utilities are, and hence it is not clear that rational people must have epistemic utilities. That is to say, I am willing to grant that rational people have all-things-considered utilities. But it is not clear to me exactly what should be 'subtracted' from 'all considerations' in order to arrive at purely 'epistemic' utilities.

Consider, for instance, my home robot servant, Hal. The robot factory equipped Hal with re-programmable degrees of belief, re-programmable utilities, a conditionalization module, and an expected utility maximization module. When I bought Hal I set his degrees of belief equal to mine, his utilities equal to mine (that is to say, my 'all-things-considered' utilities), and I instructed Hal to act on my behalf when I was not present. Occasionally Hal and I updated each other on the evidence that each of us received since our last update, and all went well. Unfortunately Hal's mechanics broke down a while ago. That is to say Hal still has degrees of belief and utilities, and can still conditionalize and compute expected utilities, but he can no longer perform any actions. He just stands there in the corner, a bit forlorn. I have not bothered updating Hal recently, since he can't do anything any more. Gibbard asks me: "Suppose you just wanted Hal's current degrees of belief to be accurate, what degrees of belief would you give him?". I answer: "I don't know. Tell me what you mean by the word 'accurate', and I will tell you what I would set them to". For instance, suppose that there is only one proposition p that Hal has degrees of belief in. Of course if I know that p is true, then I will judge Hal's degrees of

belief the more accurate the higher Hal's degree of belief in p is. That much presumably follows from the meaning of the word "accurate". But this by itself does not determine what I take to be the accuracy of Hal's degrees of belief when I am uncertain as to whether p is true or not. Nor does it even allow me to figure out the expected accuracy of Hal's degrees of belief. In order to be able to calculate such expected accuracies, I need to attach numerical values to the accuracy of degree of belief distribution/world pairs (where these numerical values are unique up to positive linear transformations.) And I don't know how to do that. So I am stuck. I suggest that this is not for lack of rationality or lack of self-knowledge on my part, but rather, because Gibbard is asking an unclear question.

Presumably Gibbard would respond that the above paragraph is confused. On his view of course the question "What would you set Hal's degrees of belief to if you just wanted them to be accurate?" does not have a person-independent, objectively correct, answer. The problem, according to Gibbard, is precisely that one could rationally have epistemic utilities such that one desires to set Hal's degrees of belief to equal one's own degrees of belief, but one's epistemic utilities could also be such that one desires to set Hal's degrees of belief to be different from one's own degrees of belief. This just goes to show that the correct answer to his question is person-dependent.

My worry, however, is that rather than that Gibbard's question is a well-defined question which has a person-dependent answer, his question is not a well-defined question. My worry is that it is a question like: "What color socks do you want Hal to wear, bearing in mind that your only goal is colorfulness?". I can't answer that question, not because I am not clear about my own desires or because I am not rational, but because the term 'colorfulness' is too vague, or ill-defined. Similarly, I worry that the term 'epistemic' is too vague, or ill-defined, so that there are no well-defined (person-dependent) numerical epistemic utilities.

4. Why have degrees of belief?

Suppose one's only concerns are epistemic. Why then have 'degrees of belief'? That is to say, when one's only goal is truth why should one's epistemic state satisfy the axioms of probability theory? I see no good reason. Let me indicate why I am skeptical by very briefly discussing standard arguments for having belief states which satisfy the axioms of probability theory.

Standard Dutch book arguments rely on the assumption that one does not want to be guaranteed to lose money, or, more generally, that one does not want to be guaranteed to lose prudential value. So, prima facie, if one's only concerns are epistemic, Dutch book arguments have no bite. However, there have been attempts to remove prudential considerations from Dutch book arguments. (See, for instance, Howson and Urbach 1989, Hellman 1997, or Christensen 1996.) The basic idea of these attempts is to claim that the epistemic states of rational people must include judgments regarding the 'fairness' of bets, where these judgments have to satisfy certain axioms which, in turn, entail the axioms of probability theory, so that, purportedly, the epistemic states of rational people must include degrees of belief which satisfy the axioms of probability theory.

There are two reasons why such arguments do not show that one's epistemic state must include degrees of belief which satisfy the axioms of probability theory when one's only goal is the pursuit of truth. In the first place the authors give no justification based only on the pursuit of truth for why epistemic states should include judgments of the 'fairness' of bets. (This may not be a slight on the cited authors, since it is not clear that they intended to give such a justification.) Secondly (and this is a slight on the authors), as argued in Maher 1997, even if a

rational person does have epistemic reasons for having such a notion of ‘fairness’ of bets, the authors’ arguments for why this notion should satisfy the suggested axioms are not convincing. In fact, Maher shows that some of the suggested axioms will typically be violated by rational people. For instance, if a person judges a bet to be fair just in case the expected utility of accepting the bet is zero, and if her utilities are non-linear in dollars, then her judgments of fairness will violate some of the proffered axioms.

The next type of arguments rely on so-called ‘representation theorems’. Such theorems show that preferences which satisfy certain axioms are always representable as those of an expected utility maximizer who has degrees of belief which satisfy the axioms of probability theory. I already find it hard to see why a rational person’s all-things-considered preferences should satisfy some of these axioms.² I find it even harder to see why a person’s purely epistemic preferences should do so, even assuming that sense can be made of ‘purely epistemic’ preferences. Let me explain in slightly more detail why I find it so hard to see why there should be purely epistemic preferences which satisfy the axioms needed for representation theorems.

One of the axioms needed for representation theorems is that preferences are transitive: if a rational person prefers A to B and B to C then she prefers A to C. When it comes to all-things-considered preferences this axiom seems to me very plausible. For, on a very plausible understanding of what all-things-considered preferences are, one can be money pumped if one violates this axiom. Now, however, let us consider the case of purely epistemic preferences. Perhaps in this case too one can be money pumped. Fine, but why should one care if one only has epistemic concerns? One might respond that the money pumping argument should not, at bottom, be taken to be a pragmatic argument which only applies to people who are concerned at avoiding a guaranteed loss of money; rather, the argument serves to demonstrate the fundamental incoherence of preferences which are not transitive. I am not moved by such a reply. It may well be that preferences can not coherently be taken to violate transitivity. However, that merely shifts the issue. For then the question becomes: is there any reason for a rational person with purely epistemic concerns to have preferences at all? I can see no such reason.

Finally, there are arguments such as ‘Cox’s theorem’, and ‘de Finetti’s theorem’, which show that ‘plausibility’ judgments which satisfy certain axioms are uniquely representable as numerical degrees of belief which satisfy the axioms of probability theory.³ Again, I can think of no non-question begging reason why the epistemic states of rational people with purely epistemic concerns should include ‘plausibility’ judgments which satisfy the axioms in question. Let me give a little bit more detail.

De Finetti’s theorem and Cox’s theorem do roughly the following: they show that one can recover the quantitative values of a probability distribution from the associated comparative qualitative probability judgments. Now, there is a way in which these theorems are not that surprising. For instance, imagine a probability distribution as represented by a heap of mud lying

² For instance, Jeffrey’s ‘continuity’ axiom and Savage’s ‘P6’ axiom seem to have no obvious justification other than mathematical expediency. See Jeffrey 1983, chapter 9, and Savage 1972, chapter 3.

³ See, for instance, Jaynes 2003, chapter 2, or Howson and Urbach 1989, chapter 3. The fundamental notions in the case of Cox are ‘plausibilities’ and ‘conditional plausibilities’, and in the case of De Finetti the fundamental notion is that of ‘comparative likelihood’.

over a continuous space. Then one can think of the qualitative probability judgments as being claims of the form: the amount of mud over area A is bigger or smaller than the amount of mud over area B. Now, clearly, one can not shift the mud around in any way without altering some such qualitative judgments. So the qualitative judgments determine the quantitative probabilities. While this argument as it stands is not precise, and does not prove exactly what de Finetti and Cox proved, it does give one some of the flavor of their theorems.

Now, while the axioms in question may seem plausible to many, this, it seems to me, is due to the fact that one has in mind that the plausibility assessments are the natural qualitative judgments associated with quantitative probabilistic assessments. One way or another, for instance, the presupposition is made that the possible epistemic states with respect to a single proposition form a 1-dimensional continuum, and no argument for this is given based on purely epistemic concerns. More generally, in so far as one thinks that the axioms on plausibility judgments can not coherently be violated by a rational person with only epistemic concerns I can see no reason why the epistemic state of a rational person with only epistemic concerns should include such judgments. So Cox's theorem and De Finetti's theorem do not seem to supply a purely epistemic justification for having degrees of belief satisfying the axioms of probability theory.

In short, I am not aware of any good purely epistemic argument for having belief states which satisfy the axioms of probability theory. Now, one might respond that, indeed, the reason for having belief states that satisfy the axioms of probability theory is (at least partly) prudential, but that, given that one has such belief states, one can ask whether rational people can have purely epistemic reasons to be dissatisfied with the degrees of belief that they have. However, if a rational person has no purely epistemic reason to have degrees of belief, why think a rational person must have purely epistemic preferences over all possible degree of belief distributions?

5. Conclusions

The notion of purely epistemic concerns is unclear to me. In so far as it is clear to me I find it hard to see a purely epistemic reason for a rational person to have belief states which satisfy the axioms of probability. If I nonetheless grant that a rational person does have such belief states and that it is clear what purely epistemic concerns are, then I can see reasons for a rational agent to have two different sets of degrees of belief: epistemic ones which serve only to maximize her epistemic utilities, and prudential ones to do everything else. Prudential degrees of belief should then be updated by conditionalization. Epistemic degrees of belief will get dragged along by the prudential ones, relegating epistemic utilities and epistemic degrees of belief to the status of an unimportant side-show.

REFERENCES

Christensen, D. (1996): "Dutch-Book Arguments Depragmatized: Epistemic Consistency for Partial Believers", *Journal of Philosophy Vol. 93 No. 9*, pp. 450-479

Hellman, G. (1997): "Bayes and Beyond", *Philosophy of Science Vol. 64*, 191-221.

Howson, C. and Urbach, P. (1989): *Scientific Reasoning, the Bayesian Approach*. La Salle, Illinois: Open Court.

Jaynes, E. T. (2003): *Probability Theory, the Logic of Science*. Cambridge: Cambridge University Press.

Jeffrey, R.C. (1983): *The Logic of Decision*. Chicago: Chicago University Press.

Maher, P. (1997): "Depragmatized Dutch book arguments", *Philosophy of Science Vol. 64*, 291-305.

Savage, L. J. (1972): *The Foundations of Statistics*. New York: Dover Publications.