

MARK NORRIS LANCE

SUBJECTIVE PROBABILITY AND ACCEPTANCE

Recently a number of philosophers have argued that Bayesian decision theory and epistemology vitiate the status of belief or, in their semi-technical terminology which I follow throughout and later explicate, “acceptance”. Some, e.g. Maher, Kaplan and DeSousa, argue that acceptance has no role to play in rational decision making while allowing that it is important to the assessment of knowledge claims and the justification of theories. Others, e.g. Jeffrey, come quite close to the more radical claim that either there is no such thing as belief or that such “digital” relations to propositions are unimportant to *any* normative enterprise, either epistemology, decision theory or ethics.¹

Kaplan puts the point as follows:

Bayesians have worked hard to show why their favored way of talking about doxastic attitudes deserves an important place in a theory of rational persons. They have done so by showing how a rational person’s degree-of-confidence function, together with her utility function, determines her decisions . . . Acceptance theorists have yet to provide a credible story about how a rational person’s acceptance of a proposition makes any difference at all to the way in which she will conduct her life . . . The point of the Bayesian argument is that, given this fact and given the indefinability of acceptance talk in terms of confidence talk, we have no account of why acceptance talk deserves any place at all in a theory of rational persons. . . .

My contention is that . . . we need degree-of-confidence talk to describe the doxastic input into rational decision-making. The proper function of acceptance talk, on the other hand, is to describe a certain feature of our behavioral repertoire – a practice in which I am engaging in the very writing of this paper.²

In this paper, I argue against such positions, attempting thereby to secure a place for acceptance talk within a Bayesian theory of rational persons. I show that Bayesian theory cannot possibly impugn the concept of acceptance, even in the limited context of rational decision theory, since Bayesian decision theory is itself conceptually dependent upon

the notion of acceptance.³ After discussing the concept of acceptance in §1, and noting my endorsement of a particular clarification of this concept which is drawn from outside the Bayesian literature, I clarify in §2 the work which I take it a theory of rational decision aims to do, thereby delimiting my target.

In §3.1, I argue that it is impossible to define the rationality of a decision strategy in Bayesian terms without presupposing a range of contingent empirical claims which are either accepted by the decision maker or assigned a subjective probability equal to one by her. Given the well known problems with the latter possibility, I assume the preferability of the former. In §3.2 I suggest a more modest version of Bayesianism which incorporates acceptance as a crucial sort of background to explicit reasoning, but which then keeps all the machinery of the traditional theory. In §3.3 I respond to three objections.

In §4, I turn to the question of the relative explanatory priority of acceptance and Bayesian probability assignments. I argue that one ought to treat Bayesian probability assignments as being in the same category as those things (propositions, assertions, sentences, or what have you) which are the object of those probability assignments. In short, I argue that Bayesianism not only must presuppose the concept of acceptance, but that Bayesian assignments of subjective probability to a proposition are best understood as themselves a special case of acceptance of a (different) proposition.⁴

1. THE CONCEPT OF ACCEPTANCE

Bayesians have adopted three different positions on what acceptance is. One is that acceptance can be identified with Bayesian degrees of confidence above a certain (nonunitary) threshold probability. Another is that acceptance may be identified with a subjective probability of one, and the final position is that acceptance must be seen as falling outside Bayesian theory altogether. On this view, if there is a role for acceptance at all, it will not be identifiable with any state definable in the language of Bayesianism.

There are many telling arguments against the first two reductive strategies and I do not rehearse them here.⁵ Given these difficulties, it seems natural for a Bayesian to accept a non-probabilistic concept of acceptance. In any event a concept of acceptance not definable in terms of subjective probability can be independently motivated and in §4 I argue that it is this concept which is required by any viable Bayesian decision theory.

Kaplan, following DeSousa, takes such an approach and defines acceptance of a proposition as follows: “My suggestion is this: that we view ‘X accepts P’ as nothing more than shorthand for ‘X would defend P were her aim to defend the truth’”.⁶ Now this notion of “defending” is not spelled out in detail by Kaplan, nor any of the other defenders of Bayesianism so far as I know. Kaplan says only that it does not imply that one would always be prepared to give a substantive argument for P, only that the agent would be prepared to assert P and to provide such defense of this assertion as is appropriate.

We can, however, offer a fleshing out of this idea as follows: there are many uses to which one might put a sentence. One is the use of standing behind it in a community of those who are seeking truth. To stand behind a sentence in this way is to take up a certain role, namely the role of the person who defends the sentence as being true in the context of argumentation. Thus, in writing this paper, I am making a number of assertions. I do not assign a subjective probability of 1 to any of them, but I do undertake to provide evidence for them and to be the target of legitimate criticism of them. I defend their truth in the context of argumentation amongst philosophers.

Acceptance, as I understand it, is largely a normatively defined state. In accepting a proposition, one undertakes a societal commitment to defend to claim in the face of legitimate challenges: to provide evidence for it, or to refute claims which are inconsistent with it. There are, of course, standards governing the range of acceptable challenges to an assertion and occasions upon which demands for evidence are appropriate. (In undertaking to defend a scientific theory, one does not undertake a commitment to answer the skeptic.)

In response to *appropriate* challenges, however, the assessor must provide reasonable justifications for her claim.⁷ There is a price associ-

ated with failure to carry out this responsibility and a reward for success. Both have to do with socially recognized *entitlement*. If argument results in someone's successfully responding to all legitimate challenges to an assertion, then that person is entitled to the assertion in the sense that they must be allowed to use it in inferences, both material and practical. (Note that on this account, if no challenges are in order in a situation, entitlement comes for free.)

Thus, entitlement to a given proposition can be used to secure entitlement either to other propositions, by providing it as evidence, or to non-linguistic acts, by showing how it makes the act rationally justifiable. Failure to adequately carry out the justificatory responsibility characteristic of acceptance results in a loss of entitlement and all that comes with it.

This account closely parallels the account of assertion developed by Robert Brandom in "Asserting". I do not know how far Kaplan or other Bayesians would follow Brandom in the details of this theory, but this is the notion that I will be concerned with in what follows, and the sense of acceptance in which I will argue that it is essential to any project of characterizing rational decision making and rational belief or assertion.⁸

In what follows, I conceive of acceptance in essentially these terms: as a state of commitment to a proposition which leads, in appropriate circumstance, to acts of asserting which are themselves socially recognized as conferring justificatory commitments – commitments to defend the proposition against legitimate criticism – upon the assertor.⁹ It is fairly clear that such a state of commitment to the defense of an assertion is an important part of our lives as rational agents. The social practice of argumentation – what Sellars has labelled 'the game of giving and asking for reasons' – is based upon this. Insofar as intellectual investigation is a social activity, it is carried on through the undertaking of commitment to assertions and their subsequent defense and challenge. Further, it is quite plausible intuitively that the rationality of a person's actions or assertions is dependent upon their commitments, upon the claims the person accepts. Change a person's commitments and you change what ought to be said concerning rationality for that person. This is the intuition which forms the backbone of the argument of §3.

2. THE CONCEPT OF A RATIONAL DECISION THEORY

My concern in this paper is with Bayesian decision theory as a *normative* theory of rational decision. That is, I am not concerned at all with Bayesianism in the sense in which this might be conceived as a descriptive theory of the way in which action is in fact produced via intentional states somehow coded in the brain. Rather, my target is a theory which purports to inform us, in some manner, about the sorts of actions an agent *ought* to perform.¹⁰

Though I take the theory to be offering normative advice, I do not assume that Bayesians purport to give us a mechanism which any agent either could, or should, *utilize* in every decision situation. The point is that Bayesian calculations may be quite complicated on any given occasion and it may well be beyond the ability of a typical agent, even a rational and intelligent one, to actually carry out the calculation, given time and other constraints. Thus, I assume only that a rational decision theory is designed to characterize, *for the theorists*, what action it is that the agent ought to perform, not how she ought to go about discovering this.¹¹

On the other hand, I do assume that a rational decision theory will serve as a potentially usable instrument *for the theorist*. The theory must provide the potential for *calculating* the correct action, given whatever information about the agent and her situation the theory takes as a starting point. Thus, though the calculation of the correct action may be so complex that it would be irrational for the agent to engage in the action of carrying out this calculation, it must be one which it is possible for *some* finite mechanism to carry out.

I also do not assume that a rational decision theory need purport to give us both necessary and sufficient conditions for an act's being rational for an agent. It is clear that this is necessary if Bayesianism, as it presently exists, is to get off the ground for utility maximization is nothing like a sufficient condition on rationality. The reason is that probability assignment to propositions may satisfy the axioms of the probability calculus and yet be irrational in ways that the agent knows about. This fact is well known and oft conceded¹² and shows that

Bayesian decision theory, like any other existing theory, merely provides necessary conditions on rationality.

Specifically, since Bayesianism has so little to say about the rationality either of valuation functions or probability assignments, which it offers is a kind of *conditional* necessary and sufficient condition on rationality. The theory tells us what it is rational for an agent to do in a situation, contingent upon the rationality of her valuations and probability assignments. This conditional task is the target of the argument of §3.1. The general form of the argument there is as follows: I consider an arbitrary specification of an agent stateable by a Bayesian solely in the language of subjective probability. I show that this is insufficient to allow one to draw any conclusions regarding what the agent should do since there will exist two agents each satisfying the given description but for whom contrary actions are rational. Thus, I argue that we need acceptance talk as well as subjective probability talk to describe the doxastic input into rational decision-making.

Finally, I also assume that a rational decision theory purports to apply systematically. Thus, such a theory must not simply tell us for some set of agents and some set of circumstances what actions it is reasonable for these agents to perform in these circumstances, given their decision theoretic psychology. Rather, I suppose that any acceptable theory must provide a principle which generates across a range of different such psychologies, a correct action for each agent having a psychology within that range.

This is just to say that we don't have a theory if we simply state that it is rational for Kaplan to drive a car; it is rational for Jeffrey to take the train and it is rational for Armendt to drink lots of water in the Summer. There must be a unity to the advice given by our theory, unity generated by the sort of principle mentioned above.

Further, if this range of applicability is restricted beyond that of all possible agents, I expect any acceptable decision theory to provide some principle which delimits this range. That is, the theory must at least tell us which range of agents it applies to and give us a principle for delimiting this range.

With these caveats in mind, we turn to the first of our central arguments.

3. WHY BAYESIAN DECISION THEORY NEEDS ACCEPTANCE

3.1. *The Main Argument*

In this section I argue that for any agent, in any decision situation in which more than one action is possible, Bayesianism can specify an act as rational for that agent only if, in its description of the relevant features of the agent’s epistemic psychology, it posits either some propositions accepted by the agent or some contingent proposition which is assigned probability 1.

Consider a typical case of Bayesian decision-making. We assume a range of possible actions $A_1, \dots, A_i, \dots, A_n$, open to the agent, a range of sentences $S_1, \dots, S_j, \dots, S_m$ describing possible states of affairs, an assessment of the utility U_{ij} of each action given a state of affairs, and an assessment of the probability P_{ij} of each state of affairs given an action. The expected utility of an action A_i , $u(A_i)$, then is

$$u(A_i) = \sum_{j=1}^m (P_{ij} \times U_{ij})$$

One ought, the Bayesian then contends, to perform that action with the highest expected utility.

Thus, in order to determine the recommended action, it is necessary to determine the utility of various outcomes given that action. Such determinations of utility can be put in the form of a table for each action under consideration. Thus:

		S_1	S_2	S_3	\dots	S_n
Matrix 1	A_i	U_{i1}	U_{i2}	U_{i3}	\dots	U_{in}

Now I emphasize, regarding this table, that the elements S_i are *sentences*. The table does not somehow list possible worlds. One could offer a philosophical theory which endorsed, in the abstract, behavior dependent upon the utility inherent in each of the possible situations which might follow an action, but in no language is it possible to give a complete description of a non-actual possible world. Thus, such a theory could not give us a concrete instance of its calculation and would be unable to give us any concrete result as to which particular act a

definable agent ought to perform in a given situation. It was assumed in the last section, however, that any acceptable theory ought to do this.

As an example of a Bayesian calculation, then, one might consider whether to go skating on the lake, given that the ice might be too thin. One would consider the two states of affairs: ‘The ice is thick enough to support us’ and ‘The ice is not thick enough to support us.’ One then determines the utility of skating and of not skating under each state of affairs as well as the probability of each state of affairs obtaining (given whether one skates, but that doesn’t matter here.) Not skating gets a rather low utility one either state of affairs, while skating produces a higher utility under the assumption that the ice is thick and a very low utility under the assumption that the ice is thin. Thus, the rational decision depends upon whether the likelihood of thin ice is low enough to balance out the excessive disutility of skating on thin ice.

		Thick ice	Thin ice
Matrix 2:	skate	+5	–20
	don’t skate	–1	–1

What is rarely remarked by Bayesian theorists, however, is that there is ice under our skates when we perform the intellectual calculation of utilities as well. Thus, we might set the utility of skating at 5 under assumption that there is thick ice and at –20 if the ice is thin. But would this be reasonable if there were snipers out in the woods who were determined to shoot anyone who skated on the ice without falling through? Certainly not. Thus, the reasonableness of our determination of utilities is conditional upon the probability of such a sniper existing out in the woods being low enough to be negligible.

The general point is that an *inference* was made from the premise that A_i is performed in situation S_j to the conclusion that utility U_{ij} results, and that this inference was defeasible. The claim that the situation and action obtain does not *entail* the conclusion that a given utility results and – the crucial point – the rationality of any action in accord with the results of this calculation depends upon the correctness of this inference.

Further, it seems plausible that this defeasibility will always be present. For any given finite information about a situation and description of an action, one can imagine information that would substantially change the rational conclusion as to the resulting utility. One can imagine, for example, that there is a dragon (or an angry god, a vampire, etc.) just around the corner who is inclined to eat you (consign you to hell, drink your blood) if one of the possible situations obtains.

It is at this point that I rely on a premise which will be assumed without argument. It is a reasonably strong premise, though one I find virtually obvious. Still, premise it is, and without it the argument of this section is substantially weakened. The premise is that for any agent, A, and for any finite description, D, of a situation in which the agent might find herself, there will be some additional sentence S such that the proper inference regarding utility for A given D is different from that for A given D&S.¹³

As I said, I do not exactly know how to argue for such a claim, but for now, if this is not plausible, I ask you to suspend disbelief for purposes of this section. At the end of §3.3 I consider what might still be correct about the argument even if this premise is false.

To return to the main line of argument, there are two ways to think of an inference from the claim that a situation and action obtain to the conclusion that a given utility results: either as enthymematic or not. If we treat it as an enthymeme, that is to say that there is at least one tacit premise involved. In the case of our decision problem regarding skating, one might suppose that the agent is presupposing that there is no sniper and that the validity of the inference to utilities rests on this assumption. But now, how do we understand this “presupposition”? The natural inclination is to suppose that it is a proposition that is accepted, but this is to concede the central claim of this section, that Bayesian decision theory cannot characterize rational decision for an agent without presupposing propositions which are accepted by that agent.

One could, however, suggest that the proposition that there is no sniper is not *accepted*, but merely assigned a high degree of confidence by the agent. In this case we need to look at a more complex table in order to perform a rational calculation, however. Thus,

	Thin ice & sniper	Thin ice & no sniper	Thick ice & sniper	Thick ice & no sniper
Skate				
Don't skate				

We need a utility for each square in this grid since it is no use to suppose that we can simply ignore the possibility of a sniper given the very low probability it is assigned. The reason is that, as we stipulated above, it is a goal of Bayesian decision theory to give us necessary and sufficient conditions for rational action contingent upon a characterization of the agent's psychology. But if we describe an agent only as assigning low probability to the possibility that there is a sniper present, then we are in no position to determine what it is rational for such an agent to do since according to the very principles of Bayesianism this depends on further facts. Some such agents' behavior will not be affected by the possibility; some will.

The problem with filling in this matrix, however, is that we cannot stop with it either, since these new utilities depend upon whether the sniper's gun is loaded. Our assumption was that *any* inference to a probability will be defeasible. Thus, it is not rational to use *any* such grid in determining expected utility unless we presuppose either a range of accepted claims or a range of empirical claims which are assigned a probability of one.

Now the argument so far was predicated on the assumption that we treat the inferences as enthymemes, but it is no more promising to treat them otherwise. To do so is to deny that there is a tacit premise in the inference and to say that one simply accepts some conditional claims along with the simple claims about the world. Thus, one simply accepts that if the ice is thin, an done were to skate then the result would be a low utility. This is not taken to be an entailment nor as shorthand for an entailment; that would not be rational for anyone. Rather, it is a conditional (a subjunctive, say) which is taken to be true in the actual world.

It is clear, however, what is wrong with this response. It has given the game away directly by talking of acceptance of a claim (albeit a conditional one). Whether one supposes that this is a matter of accepting a conditional claim or of assigning a unitary probability to the conditional, the thesis of this section holds since, by hypothesis, the conditional is not a necessary one.

It does no good to object here that one might simply *infer from* $A_i \& S_j$ to U_{ij} without ever having had any attitude at all toward the *conditional*. It is true that one might (and that one does so all the time), but so suppose that this is an objection is to conflate acceptance and belief. Acceptance is not a conscious state. All that is required for one to have accepted a claim is that one be committed to defending that claim if it is appropriately criticized in the course of rational debate.

But to take some *to be rational* in inferring from $A_i \& S_j$ to U_{ij} in the course of rational deliberation regarding action, is precisely to take them to have accepted the conditional in this sense. For any such agent is surely one who could legitimately be challenged if we were to come up with good reason to suppose that this inference was not a good one. If we were to find a reason why this conditional was false and were to present it to our hypothetical agent, this would clearly put that agent in a position from which she must, on pain of irrationality, either defend the conditional against our charge or alter her inclination to infer according to it.

Again, one could reject the idea that the conditional is accepted by supposing that the attitude is merely to assign it a probability very near to one, but this is no more promising than in the last case for now we need to consider a utility for each of the following grid boxes:

	Thin ice & $A_i \& S_j > U_{ij}$	Thin ice & $\sim(A_i \& S_j > U_{ij})$	Thick ice & $A_i \& S_j > U_{ij}$	Thick ice & $\sim(A_i \& S_j > U_{ij})$
Skate				
Don't skate				

As before, only by working out the expected utility sum which results will we know that the very low probability that the conditional is false is not counteracted by a very high utility or disutility, i.e. that our hypothetical agent is one the rationality of whose behavior is determined by the simplex matrix. Also as before, however, we cannot rest with this or any other partitioning of the space of possible states of affairs.

Thus, we conclude that rational Bayesian calculation is impossible without the presupposition that there is a range of contingent claims which are either accepted or assigned a subjective probability of one. Only in one of these ways can an agent's epistemic state justify the theorist in taking some well-defined Bayesian calculation as determinative of rationality for that agent.

Now the argument so far might seem unfair since I have been presenting the case as if the agent were using the Bayesian theory as a mechanism, something I admitted it should not be required to function as. But this was merely convenient for exposition. We can sum up our argument as follows: Bayesianism sought to begin with a set of normative psychological facts about an agent and then, largely bracketing the question of the rationality of these facts, it strives to define necessary and sufficient conditions for rationality of an action on the basis of this psychology using an expected utility calculation. So our theory attempts to define a class of agents as those who assign 0.9 probability to thick ice and 0.1 to thin ice, a utility of 5 to skating on thick ice, -20 to skating on thin ice and, hence, falling into the lake, and -1 to not skating at all. The theory then says that it is rational for any such person to skate since the expected utility of skating is $(0.9 \times 5) + (0.1 \times 20) = 4.5 - 2 = 2.5$ while that for not skating is -1 .

We ask the following of the theory: are the agents in question defined as someone who either accepts or assigns probability 1 to the sentence 'if I skate and the ice is thick then my utility will be 5' where this conditional is meant non-enthymematically? If the answer is 'yes', then Bayesianism presupposes acceptance in the process of defining agents. If the answer is 'no' then there are two possibilities: (1) the agent assigns high probability to the conditional or (2) she has no attitude at all to the non-enthymematic conditional.¹⁴

If option (1) is taken, then the agent has not been specified in sufficient detail to justify the claim that the given expected utility calculation is determinative of rationality for her. Some agents with the given psychology are rational in so inferring, others not. Attempts to specify more detail, however, lead to a regress. If option (2) is taken, then she must treat the conditional as an enthymeme. In this case there are again two possibilities: (2a) she either accepts or assigns probability 1 to a missing premise (there is no sniper), or (2b) she assigns a high probability to the missing premise. If (2a) is taken, the conclusion has been granted, but (2b) leads again to a regress.

In sum, the argument of this section shows that not enough information is included in the acceptance-eschewing Bayesian psychological description to entail the appropriateness of a given expected utility calculation. Given any such psychology, we can find two agents to whom different decision matrices apply.

3.2. *Contextually Embedded Bayesianism*

The rejection of acceptance need not be seen as an essential part of the Bayesian explanatory strategy, for we can imagine another sort of Bayesianism which finds a home *within* the practice of undertaking commitment and defending entitlement to assertions. Indeed, the previous argument should not be taken to be an objection to Bayesian decision theory per se, for the simple reason that *any* explicit decision theory – any theory which considers explicit psychological attitudes and moves from these to rationality of an action – will be subject to a similar argument showing that it must rely on a broad, unarticulated, and never completely articulable range of tacitly accepted claims.¹⁵ The only priority given to acceptance over subjective probability assignments here is that these underlying claims must be accepted. But given them, Bayesian calculations are marvellously useful.

The picture of rational decision on the basic of subjective probability assignments which begins to emerge is the following: in doing any calculation to determine the rationality of an action, one works from within a certain framework. This can be thought of metaphorically as one's world view. This world view involves a vast number of tacit and

uncritically accepted claims which define that range of relevant alternative actions and the range of possible states of affairs worth considering. Given all of this one can define a partitioning of the space of possible situations and the space of possible actions for the purposes of Bayesian calculation. But to partition the entire space of logical possibility from scratch, without the benefit of a structure of presupposition is impossible. Or to put the point the other way around, to take any specifiable matrix to determine the relevant input to an expected utility calculation which is normatively authoritative for an agent *just is* to take that agent to be committed to a huge range of claims, namely, any the explicit denial of which would lead to the matrix's being inappropriate for that agent.¹⁶

Now the postulation of accepted propositions alone will not secure for us a solid context for rationality calculation. Mere acceptance that a situation does not obtain does not imply the situation's irrelevance to decision making. The reason is that one might well accept a claim to which one assigns a low probability. If one accepts P, but assigns P a probability below 0.5 say, then this claim ought to be included in the decision matrix.

This point, important though it is, does not constitute a response to the argument for acceptance. That argument shows that Bayesianism must postulate a relation between agent and proposition which is sufficient to render it rational for the agent to leave that proposition out of her Bayesian matrix. I claim further, that this postulation must *entail* that the agent accepts certain claims. Whether acceptance alone is *sufficient* is irrelevant to this point.

In response to the positive account just sketched, however, this point shows that one must take the set of accepted propositions which form an agent's world-view, to be a proper subset of the claims that she accepts. If we want to distinguish this subset, we could label it the set of "provisionally unquestionable" propositions. These are the claims which, as Wittgenstein urged in *On Certainty*, it would be incoherent for us to claim to doubt in ordinary situations.

Unlike claims which we accept but consider fodder for serious debate, we do not, in typical situations, attend to these provisionally unquestionable claims at all. Indeed, it is definitive of their special

role as the structure inducing background against which explicit debate occurs that they be claims which ought not be explicitly considered. We may not even know we accept them until someone claims not to believe them and even then we are likely to doubt that the erstwhile challenger is experiencing a case of real doubt as opposed to merely feigned philosophical doubt.

This is not to say that the provisionally unquestionable claims are altogether unquestionable. One could easily imagine a situation in which I might doubt that there is a computer in front of me or that my street is free of snipers. Thus, sentences can come to lose this status, but they don't often do so and their present status is of a set of unargued for assumptions which we simply ought never think to doubt.¹⁷

Now I discuss these claims in some detail in the epistemology section of *The Grammar of Meaning*, but I hasten to add that I have little of detail to say about when it is rational to revise these claims or to assign specific claims such a status. However, this is no objection to "contextually embedded Bayesianism", for we stipulated up front that Bayesianism was not going to tell us anything much about the rationality of our background psychology, merely how to get from this to action which is rational contingent upon it. Thus, contextually embedded Bayesianism does just as much work as its context-free competitor.

3.3. *Some Objections*¹⁸

Before turning to the other main argument of this paper, I pause to consider three objections which one might put forward against the previous. The first, let us call "the idealization objection". According to this objection, what is going on in a Bayesian decision situation is not that the agent is accepting various claims, such as that there is no sniper. Rather, the agent realizes that the simple matrix which ignores the sniper possibility is one which oversimplifies the situation, but given the extreme unlikelihood of the existence of the sniper, she uses the simplified matrix anyway, confident that such use will not get her into any trouble.

The idea is analogous¹⁹ to the role of simplifying assumptions in Newton's calculations of the motions of planets. It is not that Newton believed that there were no additional forces acting on the planets, but

merely that these would have no substantial effect on his calculations. Similarly a Bayesian can ignore various possibilities without believing that they do not obtain.

The objection misses the point, however. Though belief is strictly irrelevant (since the issue is one of acceptance) note that Newton in the above account is taken to *believe* that there would be no substantial effect on his calculations if he were to include other force. It is analogous for the Bayesian decision agent. It may be possible for her to believe that there is no point in including various and sundry possibilities in the matrix without her accepting that these possibilities don't obtain. But this is just another way of filling in the psychology of the agent in more detail *by including accepted claims*. We distinguish the agent who has the given degrees of belief and accepts that this matrix includes all that makes a practical difference from the agent who does not accept this.

To put the point the other way around again, suppose that someone explicitly calculates according to a given matrix. You then come along and *deny*, with some plausible sounding grounds, that this matrix is sufficiently detailed to justify the rationality of an action. The agent refuses to endorse or defend the sufficiency of her matrix in argument – refuses to give reasons why it is sufficient – but insists on continuing to use it. This surely constitutes irrationality. Thus, far from being a refutation, the idealization objection actually points out another accepted claim that we must posit in order to define the range of agents for whom any Bayesian calculation is rational.

Now it is, I suppose, possible for the Bayesian to claim that the acceptance of the adequacy of a matrix as an idealization could be in the theorist's mouth rather than the agent's. That is, the theorist could say that action A is rational for any person to whom matrix M *applies*. Then, she could say that a matrix applies to an agent if the agent's assignments of probabilities and utilities to descriptions of states of affairs are such as to make the matrix a good approximation of ideally rational decision making for her. The point is to suppose that the agent has assignments of probabilities to all the various propositions which could be relevant to the applicability of the matrix and that these are such as to make these irrelevant, but that the agent neither accepts that they are such nor uses

them in decision making. Rather, we the theorists simply specify that agent we are talking about as one who has such a set of assignments.

One could say this, but even if there is such a thing as *every* proposition which could be relevant to the applicability of the matrix and even if we suppose that people have probability assignments to each of these, we must assume that these assignments will differ from agent to agent. So how do we theorists know, upon coming up to a real live agent, whether to apply a given matrix?

No finite amount of information about the agent will tell us, and the theory now gives most unhelpful advice. It specifies a small bit of psychology *S* (that which goes into the matrix *M*) and then says “apply matrix *M* to the those members of the class of *S*-agents whose psychology is such as to make *M* a good idealization of rational behavior calculation for them.” This is not a theory. Specifically, it runs afoul of our assumption that a decent decision theory ought to explain the demarcation of its ranges of applicability.

The second objection involves the assumption that an inference is made from the specification of an action and a situation to the conclusion of a utility. Rather than supposing that we infer from ‘I perform action *A* and the situation is *S*’ to ‘The utility will be *u*’ we might simply see what went on as an *assigning* of a utility to *A*&*S*. In considering this suggestion, we must first ask what exactly such an assigning is supposed to be. To assign a utility of, say, \$100 to ‘*A*&*S*’ seems on the surface to be the same as judging that the outcome of doing *A* in *S* is equally valuable to being given \$100. And I cannot actually imagine any explication of the notion which would not entail a commitment on the part of the agent to such a conclusion.

There is, however, another possible reading of the suggestion. The idea might be that the assigning is, just like the probability function, something which serves as a background to the theory. The theory is offering normative advice conditional not just on the unquestioned assumption that the agent has certain subjective probabilities, but also conditional upon her having assigned a certain utility. (Note that this assigning of utilities is, on this suggestion, not a states to affairs, but to descriptions.) We start with the agent defined as assigning *u* to ‘I skate

and the ice is thin'. Then the theory merely says, if these are rational then so is acting thus and so.

Well this can be done, but two points need to be noted. First it is quite implausible to suppose of anyone that they would engage in such assignments independently of their attitudes towards background considerations of the sort I have been discussing. This is, of course, not to say that the background considerations are conscious, but suppose a person assigns high utility to 'I skate and the ice is thick'. Would we not then expect them to deny, assign low probability, or whatever, to the suggestion that there is a sniper in the woods, *if the issue were to be raised?*

Second, even if this consideration can be ignored for *some* theoretical purpose, it is certainly the job of the theory of rational action to consider it at some point. Thus, while it might be possible to discuss what a person ought-rationally-to-do-given-their-utility-and-probability-assignments, a person who took it that there was quite likely to be a sniper out in the woods, yet assigned a high utility to skating on thick ice and who then acted on this in the usual sense would be acting irrationally. An action by this person which was endorsed by a Bayesian calculation relying on such a high probability would be an *irrational action*, even if it was *rational-given-the-utility*.

Further, we can say precisely why. It is so because, unlike us, the person does not simply accept as a contextual presupposition that there is no such sniper, or that the probabilities are such that they can ignore the possibility, or whatever. Now that is a true claim and it is part of the theory of rational action. Thus, the point still stands that we do have an account of why "acceptance talk deserves any place at all in the theory of rational persons". As was pointed out above (n. 10), it is uncontroversially true that some things can be said about rational behavior without mentioning acceptance. What is at issue is whether acceptance has an interesting role in a relatively complete theory.

The final objection I want to consider is the "denying the premise" argument. As was pointed out above, the argument for the necessity of acceptance in Bayesian theory relies on the premise that any inference from a description of a situation to a utility judgment regarding that situation would be defeasible, regardless of how much detail was

imported into the description. If one denies this, then the argument fails. That is, without the premise, I cannot demonstrate the Bayesianism must postulate acceptances in order to characterized the function from psychological state to rational action.

However, the argument is still not without interest. Suppose that there was a way to block the regress without appeal to accepted propositions. Thus, suppose that for any agent there exists a matrix so complex that the various situation description in it actually *entail* valuations of those situations for the agent. Now such a matrix would undoubtedly be enormously complex, but since we took the job of decision theory to be only to theoretically characterized rational action rather than to provide a mechanism which should always be followed, this is not to the point.

Nonetheless, though providing a useful mechanism is not a *requirement* of Bayesian theory, it would be nice. That is, if we had two accounts one of which yielded rational actions on the basic of calculations which never were nor should be carried out and the other of which yielded similar results, but in a way much closer to actual reasoning practice, we should prefer the latter.

Now my own inclinations in the direction of positive theory are those described in §3.2. That is, I think any rational evaluation of actions (or propositions) and any theoretical characterization of such, must take place within the context of an implicit posited structure of provisionally unquestioned presuppositions. And even if one thinks that it is possible to define a decision matrix which avoids the arguments of §3.1, isn't it enormously plausible that we do actually reason with simple Bayesian calculations within such a background rather than using enormous and complicated calculations without any background presuppositions?

At the very least, I suggest that there is no telling evidence against this view. Hence, acceptance-free Bayesianism at least has a competitor in "context embedded Bayesianism." I think this latter has intuitive plausibility on its side, but to argue for it in detail would require a much longer paper. At the very least, though, I take the argument of §3.1 to show that Bayesianism as it is conceived by the opponents of acceptance has nothing to show for itself over this alternative which embraces acceptance. (It may be thought that simple economy of posits

was a virtue, but in the next section I show that Bayesianism involves acceptance at another point anyway.)

4. THE CONTENT OF SUBJECTIVE PROBABILITY ASSIGNMENTS

Depending upon whether the reader accepts the crucial premise of §3, and assuming she rejects assigning unitary probabilities to contingent, empirical claims, it has been argued either that Bayesianism *must* posit a background context of accepted propositions or that there is an alternative theory which makes use of any Bayesian mechanism which an actual agent could utilize in real situations, and which posits accepted propositions. In this section, it is argued that the very subjective probability assignments of Bayesianism themselves should be understood as accepted propositions. The view I will argue for is that to assign subjective probability n to a proposition P , *in the sense in which such assignments are relevant to the normative project of rational decision theory or epistemology*, just is to accept the proposition that fair rational betting odds on P are $1-n/n$. (In the interest of brevity I will simply formulate this as “the rational subjective probability of P is n ”.)

Let us begin by recalling the fact that Bayesianism purports to be a normative rather than a descriptive theory. Unfortunately this distinction is not always carefully attended to in the presentation of Bayesians. Thus, some suggest that we can equate subjective probability assignments with dispositions to bet. Even Jeffrey (in an uncharacteristically cryptic remark) calls the usual matching of betting dispositions with subjective probability assignments a “quasi-operational definition”.²⁰

It is clear, however, and generally accepted, that this is a mistake if the goal is to explicate a notion of rationality. It is possible to be disposed to bet at certain odds and to act accordingly, but to be irrational nonetheless as is shown by, for example, the existence of compulsive gamblers.

This is widely recognized by Bayesians and few would claim that the rationality of a degree of belief assignment is given simply by the formal constraints placed on it by the laws of the probability calculus. Thus, Jeffrey has emphasized that his ignoring other constraints is not because he believes that satisfying the probability calculus is sufficient for the

rationality of a probability assignment or the behavior based upon it, but simply because there are no other precisely formulable necessary conditions available at present.²¹

The problem may be deeper than present ignorance, however. I am not confident that there *are* any conditions on the rationality of subjective probability judgments, other than Bayesian ones, which are universally valid and formulable with mathematical precision. Even if there are, I am quite confident that no one will ever present necessary and sufficient conditions for rationality from this abstract a perspective. Rationality is always a matter that is criticizable on the basis of empirical knowledge, the ultimate reason being that there is no more a clear or stable line between rationality and knowledge than there is between the analytic and the synthetic.

Thus, one can expect of a theory that it make room for substantive empirical criticism of degree of confidence assignments. The theory may well not tell us anything about how such substantive criticism works, but it must be consistent with the intelligibility of empirical criticism of probability assignments.

To take a concrete example, suppose that George has seen a coin flipped 100 000 times and seen it land heads each time. It may still be perfectly coherent, in the technical Bayesian sense, for George to feel that the probability of heads is 0.5. There is no reason why such a degree of confidence need be inconsistent with any of his other assignments.

However, it is essential to our practice of evaluating each other's epistemic behavior that we take George to be irrational. The reason is that he has offered a hypothesis which flies in the face of the evidence. Any rational person would conclude that the coin is nearly certain to land heads every time and would appropriately demand of George that he explain why he disputes this.

If George bets at even odds that the coin will land tails, we will ask him why he would do this. He, being a good Bayesian, might say that his subjective probability assignment to the proposition that it will land tails next time is 0.5 and that this is perfectly coherent (in the technical sense). We will argue with him, however, insisting that this is foolishness, that fair betting odds are nowhere near even, that the evidence clearly suggests that the right probability is near 1 for heads,

etc. Thus, disputes about the rationality of action quickly give way to ordinary disputes about the rationality of propositions regarding fair betting odds.

Now there are two possibilities here. Either this could be explained on the basis of the assumption that what action is to be judged against simply is one's set of accepted propositions regarding fair betting odds. Alternatively, one might insist that behavior is to be evaluated against assumptions about one's mere assignment of betting odds while allowing that there can then be debate – debate involving propositional acceptance – about the rationality of this assignment.²²

There is reason to prefer the former course, and it has to do with the ease that understanding of the situation has in handling a range of phenomena. There are several other categorical features of propositions which are exemplified by the contents of those psychological states on which conditional rationality of action depends. First of all, the states iterate. Consider the following case: Jane is not be sure which of two pennies she purchases from the local crooked gambler. In such a situation, Jane might take the rational subjective probability to be 0.4 for the proposition that the probability of this coin landing heads when flipped is 0.5.

Such nested judgments have consequences for the rationality of first order judgements. If Jane is as above and takes there to be two possible coins she might have been given, one fair and the other weighted so there is a 0.9 probability of heads, we can calculate as follows. She judges that the rational probability is 0.4 that the rational probability of heads is 0.5 and that the rational probability is 0.6 that the rational probability of heads is 0.9. Thus, if she were to bet on the overall chance of the coin landing heads with this background, she should take it to be $(0.4 \times 0.5) + (0.6 \times 0.9) = 0.2 + 0.54 = 0.74$. She ought in such a situation to set fair betting odds at roughly 3/1.

Now this way of discussing the calculation is quite natural and takes the psychological states in question to be accepted propositions. Jane accepts that fair betting odds on the proposition that fair betting odds are such and such are so and so. This is not the only way to understand the situation, however. One could still insist that subjective probability assignments are merely assignments, not acceptings of propositions. As

before, one cannot take this to be all that is present psychologically, for if assignments are not themselves propositions, they cannot be within the scope of a probability assignment. Probability functions or assignments do not have likelihoods, only propositions that the function is thus and so do.

As in the discussion of rational debate regarding probability functions, however, the Bayesian could insist that what we have are both assignments and accepted propositions about those assignments. The task, then, will be to motivate the above inference. Mere demands of probabilistic coherence do not require any connection between one's subjective probability regarding the likelihood of a coin landing heads and one's subjective probability regarding the proposition that the rational likelihood of the coin landing heads is 0.5. But decision theory ought to recognize such a connection, so the rules of coherence should be augmented with inferential rules governing such connections. This could, however, be done.

In addition to nesting, one can make inferences from commitments to defend the fairness of betting odds. One can also express these inferences explicitly as various sorts of conditional. In the story of George, for example, one might correctly say: "*If* the probability of the coin landing heads is 0.5, as you assumed in agreeing to bet at even odds, *then* we have no particular reason to assign high probability to the sun rising tomorrow since we have less evidence about the constancy of that than we do about the constancy of this coin."

Once again it is easy to understand this with the assumption that George begins simply with an accepted proposition to the effect that the rational subjective probability of the coin's landing heads is 0.5. Also as before, however, it is possible to maintain that what we have here is both a non-propositional assignment of probability to a proposition together with an additional propositional attitude about that assignment. It is then this latter which figures into conditional judgments and which is connected by bridge principles to the assignments of probabilities, and from there to actions.

Now we must ask, though, why all this complication? Why have two sorts of states at all? If the rationality of actions is to be considered conditional upon a set of psychological states which are such as to carry

with them commitments to defend propositions when behavior is challenged, which are logically tied to nestable propositional contents, and which are logically tied to contents which figure in inference and conditionals, why not simply accept that the states in question are themselves propositional?

If this simplifying move is to be convincing, more must be said about what *sort* of judgment this input into decision theory is. According to the account of acceptance sketched above, to claim that fair betting odds on proposition P are 1-n to n is to undertake a responsibility to justify this claim in the face of reasonable challenges. Successfully carrying out such justification (which, again, could be achieved vacuously in the absence of reasonable challenges) carries with it entitlement to the proposition accepted. This entitlement consists in the permission to use the proposition in inference, both material and formal inference to other claims, and practical inferences to actions.

This latter role for entitlement – that of justifying actions on the basis of claims to which one has secured entitlement – allows us to incorporate the entire Bayesian theory of rational action into the present framework. The reason is that the fundamental practical inference which ought to be associated with such propositions – propositions that fair betting odds on P are 1-n to n – is that if one is entitled to this judgment, she is thereby rationally justified in placing bets at any odds at least this favorable. To say that she is so justified is to say that one cannot deny the rationality of the betting action without challenging the agent's entitlement to the judgment.

This theoretical context for Bayesian degrees of confidence assigns them an explicitly normative role. The point of justifying propositions regarding the subjective probability of other propositions is typically to justify actions. One is justified in acting insofar as these actions are in accord – in the sense defined by Bayesian decision theory – with one's justifiable estimates of subjective probability. So the reductive *equation* of betting dispositions and assignments of probability is replaced by a normative inferential propriety. One is entitled to bet on P at odds of better than $(1-n)/n$ insofar as one is entitled to assign probability n to P. Thus, if it is pointed out that one's betting dispositions do not match

one's probability assignments, one is required to change one or the other on pain of irrationality.

Now we must not suppose that in judging that the rational probability of P is n , for example, one is committing herself to a claim concerning the *objective* probability of P obtaining. Much of the philosophical underpinnings of Bayesianism rests on this very distinction. What is relevant to the rationality of one's actions is how reasonable, given one's evidence, various propositions are, not how probable their truth is objectively. Thus, for example, one can be quite sure that whether it will rain tomorrow in Syracuse is fully determined by the present state of the world and, hence, that the objective probability of rain is either 1 or 0 (probably 1) without thereby holding the subjective probability of rain to be either extreme value. As long as one is rationally unsure, given the evidence, whether it will rain, one's subjective probability assignment to the proposition that it will is going to be strictly between 0 and 1.

Some Bayesians have taken the fact that objective probability is not what is at issue to be an argument against the coherence of basing action on judgments of probability. Maher, for example, simply moves from the claim that the content of a belief does not concern objective probability to the supposition that it concerns merely one's own psychological states. In discussing the proposal that scientists might believe sentences like 'We are rationally entitled to believe that many non-trivial empirical propositions are probably true', Maher says the following: "Presumably these probabilities are not intended to be subjective probabilities, since that would reduce a scientist's belief in a theory to a belief about himself. But on the other hand, it is difficult to understand what it could mean to say that a scientific law has a certain objective probability of being true."²³

Maher seems to be relying on a certain alleged dichotomy. On the one hand there is objective probability. Objective probability is the probability of radioactive decay on the usual interpretation of quantum mechanics, a feature of nature quite independent of any fact about people, their psychology, their history, or their society. In particular, it has nothing to do with our epistemic condition; one who knows every fact about the world up to a time would know only the probability of an atom's decaying in the next minute.

On the other side of the dichotomy is a purely descriptive or dispositional notion. To believe that the subjective probability of P is n , on this understanding, is to believe something concerning one's own psychological states. Such a subjective probability is subjective in the strictest sense of being whatever the individual makes it. Let us call this notion of probability, since it is purely descriptive, the "dispositional probability" of an act.

This is, however, a false dichotomy. One can make judgments about probabilities which are neither of these sorts of judgments. To see this, consider an analogous case: that concerning the legal facts. Few people (outside of recent appointees to the supreme court) are inclined to be objectivist about the law. There are no legal restrictions out there in nature independently of human interaction. On the other hand, no one is going to be a subjective dispositionalist about the law either, supposing that in making a legal claim, one is merely describing one's psychological attitudes. If this were their content, there would be little room for debate (at least among the non-self-deceived). One coherent feeling as to what the law is would be as good as any other.

Rather, legal discourse inhabits an intermediate *social normative* ground. The law exists within the context of an elaborate social practice. There are social standards for passing laws and then, when specific legal issues arise, there are evidential criteria determining the appropriateness of various ways of justifying hypotheses. One must refer to existing statutes and argue for a connection between one's judgment of the particular issue and the explicit claims of the statute. There are also any number of different strategies for defending various extensions, interpretations and revisions of the statutes.

The same position is possible in epistemology and, derivatively, decision theory. To say that the rational probability of P (given a range of data) is n is to perform an act that carries with it a socially recognized justificatory burden. One must argue that the data supports the claim to a certain extent, bringing in such criteria as empirical adequacy, explanatory coherence, simplicity and the other usual suspects. This practice need not presuppose that there are evidential relations out there in nature apart from the interests, biases, historical and social circumstances, and psychological peculiarities of people. At the same time, however, it

need not suppose that these judgments are not amenable to rational criticism.²⁴

So we have three sorts of probability: had I been the original name-giver I would have called them 'subjective', 'epistemic' and 'objective' since the former is whatever one makes of it – in the decision theory case one's disposition to bet at certain odds – the second is what is relevant to the rationality of action and assertion – one's normative judgment that fair odds are thus and so – and the latter has to do with the way things are in themselves – the objective probability that P is true. Unfortunately Bayesians have already used the name 'subjective' probability for the second of these notions. (That is, whether they have conceived of subjective probability in the propositional terms being developed here or not, they have used this term to name the sort of probability which must shoulder a burden only bearable by the epistemic concept.) Thus, I stick with 'dispositional probability', 'subjective probability' and 'objective probability'. The first and last are purely descriptive – the first of a psychological state and the latter of things in the world – while the second is normative.

A fuller account along the lines sketched here, would give more detail as to the inferential role of 'fair betting odds for P are 1-n to n'; it would spell out material inferences typically appropriate both to and from assertions of this form. There are many types of reasons one might present for believing such a claim and although specifying them with precision is not easy – specifying *all* of them with precision is impossible – they are not hard to recognize in practice. We do so virtually every time we argue with a person about the rationality of a bet they are inclined to make. It should now be clear that the major part of the task of making this role explicit converges with that of producing an adequate theory of justification.

All we need to notice at present is that there is a clear difference in inferential role between judgments concerning objective probability and those of subjective probability. In order to justify a claim concerning objective probability, one needs to argue for a particular value for the likelihood of P given all that is *in fact* the case, whether we know it or not. Thus, to show that some future tense P has a non-zero and non-unitary objective probability is to argue that no matter how we might

fill in our ignorance concerning the present facts, the truth of P would still be up in the air.

In order to argue that the epistemically rational subjective probability of P is n, one need argue only that the evidence presently available to us supports P to degree n, that *given our limited evidence* it makes sense for us to bet on P in the usual way. (In order to justifiably claim merely that one's dispositional probability assignment to P is n, one need only convince us that she is confident of P to degree n.)²⁵

The present account of rational action, as based upon accepted propositions concerning rational probability assignments, also provides a natural understanding of the Bayesian requirement of coherence. What precisely is the requirement that one's assessments of probabilities satisfy the probability calculus? It is certainly not, as is now widely recognized by Bayesians, that one is irrational if she violates this injunction since there are thoroughly common cases in which the violation is not something that one could rationally be expected to be aware of. Indeed, there are cases in which one is rationally required *not to* have a probability assignment which accords with the probability calculus.²⁶

Much more plausible, I submit, is that the point of Bayesianism's requirement of coherence is that one's assessments of epistemic probability *can be legitimately criticized by pointing out* that they violate the axioms of the probability calculus.²⁷ Whenever we make any claim, one of the commitments we undertake is to defend it against logically based challenges. Thus, such axioms as: If A entails B then $\rho(B) \geq \rho(A)$ can be taken as *inference license*. They are licenses upon the presentation of which, a challenger is permitted to criticize one who accepts that A entails B merely by pointing out that she does not hold that the epistemic probability of B is at least as great as that of A.²⁸ Thus, the principles of the Bayesian account of coherence – the probability calculus – are seen as partially determining the formal inferential role of judgments concerning epistemic probability.

There are concrete advantages to a Bayesianism which is explicit in its recognition that "subjective probability assignments" are themselves judgments concerning fair betting odds and such a recognition is implicit at any rate, in the Bayesian recognition that subjective probability assignments can be rationally criticized. In addition, as was shown in

§3.1, even these probabilistic judgments can themselves only do work if they function in conjunction with accepted non-probabilistic claims. So Bayesianism is doubly in need of acceptance. Bayesians, we might say, by committing themselves to their overall explanatory framework, have accepted acceptance – whether they are aware of it or not.

I would like to thank Jonathan Bennett, Patrick Maher, Mark Kaplan, and Brad Armendt for their many helpful comments and objections to earlier versions of this paper.

NOTES

¹ See Patrick Maher: “The irrelevance of belief to rational action,” Mark Kaplan: “A Bayesian theory of rational acceptance,” R.B. DeSouse: “How to give a piece of your mind: or, the logic of belief and assent,” Richard Jeffrey: “Valuation and scientific hypotheses,” and “Probable knowledge,”

² A Bayesian Theory of Rational Acceptance,” pp. 310, 311.

³ Though the argument of this paper is directed toward Bayesian decision theory, I believe that similar considerations apply to a much broader range of theories. Not only other sorts of theories which base themselves on degrees of belief, but any explicit theory taking us from psychological states to prescriptions regarding the rationality of action ought to be subject to a similar argument. In part I focus on Bayesianism because I take it to have been among the most successful such theories.

⁴ Although there are a number of influences relevant to my coming to think about this approach to degree of confidence assignments, one particularly salient one was a remark by Peter Van Inwagen that this seemed a natural approach. I am grateful to him for the suggestion.

⁵ See, e.g. Kaplan, “A Bayesian theory of rational acceptance”.

⁶ “A Bayesian theory of rational acceptance” p. 311.

⁷ It is worth noting, though, that these justifications can involve deference to appropriate experts. Even in this case, however, it is crucial to keep clear the distinction between asserting that an expert has an adequate justification for a claim – something one might claim without actually have accepted the claim oneself – and actually undertaking a commitment to defend the claim oneself by way of an appropriate gesture to the epistemic labor of another.

⁸ Even if this notion could be argued to be substantially different than what was intended by the Bayesians I am considering, it is still a thoroughly binary relation to a proposition, and hence more in the belief camp than the degree-of-belief camp.

⁹ It is interesting to note that ‘commitment’ is ambiguous in English. On the one hand it has a purely normative meaning as in ‘Jones, by virtue of having signed the contract, was committed to delivering the package by 6:00 a.m. even though he would just as

soon have chewed his own arm off.’ On the other, it has a dispositional meaning as in ‘Jones was deeply committed to living in as filthy a condition as the law would allow.’ Although I avoid, in this paper, the knotty issue of the analysis of the ordinary notion of belief, it is not unreasonable to understand a belief that P as a standing (perhaps self-aware) disposition – a commitment in the causal sense – to perform an act which commits one, in the normative sense, to P. Thus, the concept of belief involves and, to an extent, runs together the two notions of commitment, between which the English word is ambiguous.

¹⁰ It seems to me, after talking to at least five prominent Bayesians, that there are at least as many conceptions of the goal of Bayesian decision theory as there are defenders of ‘it’. Indeed, in a single afternoon I heard three of these insist that it was quite clear that Bayesianism was the doctrine that such and so. Unfortunately no two of the three doctrines presented were jointly consistent.

With this in mind, I make to claim to be criticizing any particular person’s version of the theory. I take it that my argument is non-trivial, since the sort of theory I argue requires acceptance is one that is undeniably important. Not only is it what is most typically called Bayesianism in texts on decision theory, but it is something that matters philosophically. Developing such a theory is something we clearly ought to strive for.

No doubt there are definable classes of true things which could be said without mentioning acceptance, but the interesting question must be whether some reasonably unified and complete explanatory project can be completed without mentioning acceptance. As Kaplan puts it in the passage quoted above, whether “acceptance talk deserves any place at all in a theory of rational persons”.

¹¹ The importance of emphasizing this aspect of my argument was made clear to me by comments kindly sent to me by Patrick Maher.

¹² See, for example, Jeffrey “Dracula meets wolfman: acceptance vs. partial belief”.

¹³ Indeed, it seems to me that given any language, and any inference from a description of a situation to a utility for an agent, where the description is formulated in the language, there will be such an additional sentence existing in some richer language the content of which is not even expressible in the original language.

¹⁴ Of course as Brad Armendt pointed out to me, the agent might also assign low probability to this particular conditional taking u_{ij} to be a sort of mixture of the values of various unspecified possibilities. But then we simply change the conditional to “if A and S obtain then the appropriate mixture of the various possible outcomes under various specifications is u.” That must either be accepted or assigned high probability, and the same argument runs through.

¹⁵ John O’Leary-Hawthorne and I argue for such a position in the chapter on normativity in *The Grammar of Meaning*.

¹⁶ The background to Bayesian calculation that I am calling a world view, need not be conceived of only as a range of accepted propositions about the world. Rather, it can be conceived of as a set of inferential dispositions (or, equivalently, a set of accepted conditionals). This corresponds to the non-enthymematic case of the argument.

¹⁷ Mark Kaplan pointed out in correspondence that my argument at this and a number of other points is quite akin to Levi's (in *Gambling with truth and The Enterprise of Knowledge*). Levi has no use for the concept of acceptance defined here at all, so far as I know. He occasionally calls the assignment of probability 1 'acceptance', but it is manifest that one can be committed to defending a claim without any such assignment. Nonetheless, there are many similarities between Levi's lines of thought and my own, and I do not wish to deny the affinity. It is certainly true that Levi also wants to find a status for a range of claims in virtue of which they function to define a sort of context within which the explicit Bayesian calculation takes place.

Perhaps the primary difference is that he takes assignments of probability 1 to do this work while I embed the Bayesian calculation within a context of acceptance, but I am not sure. I suspect that a detailed comparison of the positions would be quite fruitful though it would have to be carried out by someone more familiar with Levi's work than myself.

¹⁸ The first objection considered in this section was offered in correspondence by both Mark Kaplan and Patrick Maher, though the formulation offered here is my own and they should not be taken to be responsible for defending either this or any other formulation of the objection. The second objection was offered by Brad Armendt, though again its formulation is my own responsibility.

¹⁹ Patrick Maher pointed this out in conversation.

²⁰ "Dracula meets wolfman . . .".

²¹ See, for example, Jeffrey: "Dracula meets wolfman: . . .".

²² I am grateful to Brad Armendt for clearing up a serious error in an earlier version of my discussion at this point.

²³ Maher, *op. cit.*

²⁴ Of course, one *could* also take the line that Carnap did, of supposing that there are objective evidential relations in the world and using these to determine the rationality of actions. Such a line faces many (to my mind insuperable) problems, but I do not argue against it here. The reason is that it is to adopt the point of this paper, which is that one should take rational decision to be determined by acceptances concerning the rationality of subjective probability judgments.

²⁵ This social account of the content of subjective probability claims is not the only possible way to provide a third way between subjective probability and objective probability. One could claim, for example, that the probability assignments relevant to decision theory are claims of the form: "dispositional probability assignments of n to P would lead to pragmatically useful behavior." The differences between this formulation and the one considered in the text are subtle, but real. Jonathan Bennett suggested this possibility to me and it is worth pursuing. I do not take myself to have argued against this propositional content version of probability assignments in the text.

²⁶ Suppose we have, in front of us, a chess position C . Let P be a statement of the rules of chess together with a description of C . Let Q be the claim that there is a winning strategy for white in C . Now suppose, given current theory, human calculational abilities, etc.

that everything points to the truth of Q. Any grandmaster would determine that Q is very likely true. So, a typical grandmaster would assign quite a high probability to P and a slightly lower one to Q, say 0.99 and 0.9 respectively.

Suppose, however, that by a maneuver which is enormously complex and goes against current theory, there really is a win for black in the position. Thus $\sim Q$ is true. Further, $\sim Q$ is entailed by P. Thus, to be consistent with the principles of the probability calculus, our grandmaster should assign $\sim Q$ at least 0.99 and Q no more than 0.01. But this is not merely not required by rationality, it would be positively irrational for someone in this situation.

²⁷ I recall mark Kaplan suggesting such an interpretation in conversation. I may have gotten the idea from him or we may have come to it independently. I really don't remember. (Of course, the present formulation is my responsibility entirely.)

²⁸ Philip Kremer and I have discussed this general strategy for interpreting inferential principles and developed the idea in some detail for one sort of entailment conditional in our papers "The Logical Structure of Linguistic Commitment I". and "The Logical Structure of Linguistic Commitment II".

REFERENCES

- Brandom, Robert. "Asserting," *Nous* 17, pp. 637–650.
- DeSousa, R. B. (1971) "How to give a piece of your mind: or, the logic of belief and assent," *Review of Metaphysics* XXV, no. 1, (Sept. 1971), pp. 52–79.
- Jeffrey, Richard (1956) "Valuation and scientific hypotheses," *Philosophy of Science* XXIII, 3 (July 1956), pp. 237–246.
- Jeffrey, Richard (1968) "Probable knowledge," in I. Lakatos, ed. *The Problem of Inductive Logic* (Amsterdam, North-Holland), pp. 166–180.
- Jeffrey, Richard (1983) *The Logic of Decision*, second edition (University of Chicago Press, Chicago).
- Jeffrey, Richard. "Dracula meets wolfman: acceptance vs. partial belief," in Swain (ed.) *Induction, Acceptance, and Rational Belief*, p. 160.
- Kaplan, Mark (1981) "A Bayesian theory of rational acceptance," *Journal of Philosophy* LXXVIII, no. 6 June 1981, pp. 305–330.
- Lance, Mark (1989) "Rules, Practices and Norms" in *Ludwig Wittgenstein: A Symposium* (Longwood Publishers Group).
- Lance, Mark and John O'Leary-Hawthorne. *The Grammar of Meaning: Reflections on the normative character of semantic discourse*, forthcoming.
- Lance, Mark and Philip Kremer (1995) "The logical structure of linguistic commitment I: four systems of non-relevant commitment entailment," forthcoming in *The Journal of Philosophical Logic*.

- Lance, Mark and Philip Kremer (1994) "The logical structure of linguistic commitment II: systems of relevant commitment entailment," *The Journal of Philosophical Logic* 23: 369–400.
- Levi, Isaac (1967) *Gambling with Truth* (Knopf, New York).
- Levi, Isaac (1980) *The Enterprise of Knowledge* (MIT Press, Cambridge, MA).
- Maher, Patrick (1986) "The irrelevance of belief to rational action," *Erkenntnis* 24 (1986), pp. 363–284.

Department of Philosophy
Georgetown University
Washington, DC 20057
USA