Introduction to A Logic of Assertions

Robin Giles

1 Two Limitations of Classical Logic

In respect of its application to statements describing beliefs about the real world the value of classical logic is limited in two respects.

First, classical logic can deal only with statements which are "sentences" in a particularly narrow sense: not only must a sentence have, under any given conditions, a *truth value*, TRUE or FALSE, but its whole meaning must be given by prescribing the conditions, *i.e.*, the "states of the world" under which it is true. Thus a classical sentence, *i.e.*, a sentence in the sense of classical logic, is in principle completely described by giving its *truth function*, a function on the set Ω of all *states of the world*, taking values in the 2-element set $\{0, 1\}$, or equivalently {FALSE, TRUE}, of truth values.

Now, it is clear that a typical natural language statement is not of this type. Take, for instance, the standard example, *John is tall.* If this were a sentence of classical logic, then it would have a truth function, f say, presumably a function of the height h of John,[1] of the form shown in Figure 1. There is some particular height, say 5'11", at which the truth value changes instantly from 0 to 1. But this means there is a dramatic difference between an agent's assertions of this sentence when John is very slightly shorter than 5'11" and when he is very slightly taller. It is obvious that in real life this is not the case: in fact, there is no point at which a very small change in height would produce a large change in the acceptability of the assertion. This shows that if the statement *John is tall* has a "truth function" f at all it must be a continuous function of the height h of John. So, since all heights are possible and heights form a continuum, the set of truth values cannot be $\{0, 1\}$ but must be a connected set. It could, for instance, be the interval $[0, 1]$ with the truth function given by a graph such as in Figure 2. The ordinate $f(h)$ corresponding to a height h is a number

[1] We assume for simplicity that John belongs to a homogeneous population of adult humans. If this is not so — if John may be a horse or a child or a it pygmy, for instance — then the truth function still exists but the truth value is no longer a function of the height of John alone.

H.E. Kyburg, Jr. et al. (eds.), Knowledge Representation and Defeasible Reasoning, 361–385.

in the interval [0, 1] which may be referred to as the *degree of truth* of the statement. A statement like this, which cannot be accepted as a classical sentence but whose meaning can perhaps be represented in this way, is called a *fuzzy sentence*. In contrast, a sentence in the sense of classical logic is referred to as a *crisp sentence*.

On reflection, one sees that the sentence "*John is tall*" is not exceptional: almost all statements occuring in natural language are unacceptable as classical sentences but could be admitted as fuzzy sentences in the above sense; some examples are: "*That was a good dinner*", "*The toast is burnt*", "*Mary is careless*". Even statements such as "*John is a Welshman*" or "*Mary was born in 1948*", that appear to be classical sentences at first sight, can be seen on closer inspection to admit "borderline cases" in which the truth value is moot.

The above argument shows clearly the inadequacy of classical logic as a tool for coping with natural language statements, and it forms the basis for the currently popular "fuzzy set theory" approach to reasoning under uncertainty. However, although the argument suggests the introduction of a continuum of truth values and in particular the use of the interval [0, 1], it provides no justification for this procedure. More important, because it assigns no *meaning* to the concept of *degree of truth* it is not possible to determine how *reasoning* should proceed in such a system. So the argument is not sufficient: it shows the need for a logic that in some way can represent a "continuum of truth values" but it falls short of providing a foundation for such a logic.

The second, and quite distinct, way in which classical logic is inadequate as a tool for handling natural language statements is that it lacks any built-in procedure for expressing *degrees of belief*. It is true that Bayesian methods (expressing beliefs by giving subjective probabilities) can be used for the expression of degrees of belief which are exact in this sense, but they cannot be used in the case of uncertainty — in situations where an agent is not willing to assign any exact probability to the sentence in question. Moreover, the Bayesian approach is simply not applicable when the beliefs in question relate to fuzzy sentences, since there is then no "event" to which the notion of probability can be applied.

The problem, then, is to construct a formal logic, and of course a corresponding formal language, that will overcome these two deficiencies of classical logic. Now, we already have an informal language in which the deficiencies are overcome — namely, natural language itself. In spite of its ambiguities and lack of clarity natural language certainly provides facilities

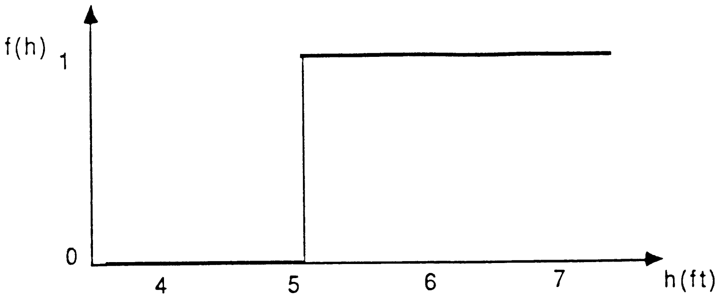

Figure 1

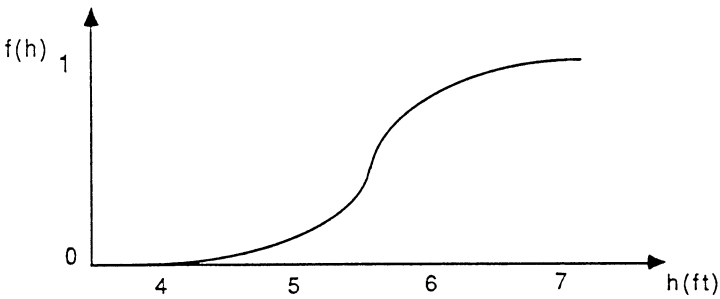

Figure 2

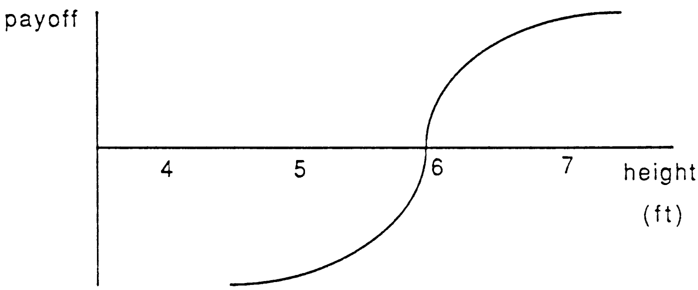

Figure 3

for formulating fuzzy sentences and expressing degrees of belief. In attacking the problem, then, it is natural that we should first try to see how natural language works in these two respects. By using it as a guide we may hope to reach a formal system that solves our problem.

2 Payoff Values in Natural Language

Let us try to find a precise way to represent the *meaning* of a natural language statement. Following the method of pragmatism [James07], we take the meaning of a statement to be determined by *the way it is used.* For example, insofar as the sentences "*It is raining*" and "*Il pleut*" are used under just the same circumstances we are justified in saying that they have the same meaning.

Now, a use of a sentence is an assertion, so we are faced with the question, How does an agent decide whether or not to assert a given sentence? Well, an assertion is an act; the agent must choose between the acts of *asserting the sentence* and *not asserting the sentence.* So the question falls within the scope of *decision theory.*

The best formulation of decision theory, particularly in the case of vagueness or uncertainty, is still to some extent under debate. The most generally accepted treatment — which we shall adopt for the present — is the Bayesian approach, based on subjective probability and expected utility. According to this view[2] any act leads to an outcome that has a certain *utility.* The outcome, however, depends not only on the act but on the current *state of the world.* As a result, the choice as to whether to make an assertion a or not is determined by the *payoff function* $a(\omega)$ of the agent where, for any world state ω,

$$a(\omega) = (\textit{payoff if assertion is made}) \; - \; (\textit{payoff if it is not made}),$$

the *payoff* being the utility to the agent of the outcome (for the given world state). Thus $a(\omega)$ can be described as the gain in utility due to making the assertion.

The payoff may usefully be described as a measure of the *willingness* of the agent to make the assertion. When positive, it constitutes his *motive* for making the assertion; when negative, it is the reason he does not make it.

[2]See, for instance, [Luce&Raiffa57].

According to this analysis it is the payoff function that determines when an assertion is made. So the payoff function can be taken as representative of the meaning of the assertion.

As an example, consider the statement "*John is tall*". To identify the payoff we ask, what is the motive that causes an agent to assert this sentence? In a few cases, for instance if the agent is John's father and John is a candidate for a basketball team, the motive may be the desire to convince the listener of the truth of the statement — if the listener is convinced then the agent profits. But usually the agent has no personal stake in the statement and the motive is simply the satisfaction arising from the approval that, in normal society, is accorded to him who makes an assertion that later proves to be justified; similarly the normal reason for *not* making a statement is the hope of avoiding the disappointment that for a normal person accompanies the loss in prestige which eventually results from the making of a misleading assertion. Thus it is not simply the form of words used but the reaction of society to the assertion that in the end is responsible for the form of the payoff function and so for the meaning of the assertion.

From this it follows that the payoff function for a given assertion depends, for normal people, only on the assertion and not on the individual asserting it. In other words, we may assume that all (normal) agents agree as regards the payoff function associated with any given assertion. This is an exact analogue of the presumption in the case of classical logic that all agents are agreed on the truth function corresponding to any (classical) sentence. In both cases the assumption represents the supposition that all agents "speak the same language", in that they all attach the same meaning to any given sentence.

Of course, to arrive at the representation of meaning by a payoff function we must assume that any agent we consider is "normal" in that he gets pleasure from society's approval and disappointment from its disapproval. Clearly, some agents are not "normal" in this sense: idiots, liars, and pranksters are examples. They get their satisfaction from other sources than the approval of society, and the payoff function for an assertion in the case of such an agent may be quite different from the norm. In the following we assume that every agent is "normal" in the sense explained.

In the case of the statement "*John is tall*" the general nature of the payoff function is easy to see: if John is definitely tall, say he measures 6'6", then the payoff is positive and large; if he is far from tall, say 4'6", then it will be large negative; and in general it is an increasing function of the height of John, see Figure 3.

We see that the graph of the payoff function is qualitatively similar to that of a "continuous truth function", assuming we interpret large-negative and large-positive as equivalent to false and true respectively.

For clarity, an explanatory remark should be made in connection with Figure 3. As was noted above, the payoff is really a function of the whole world state. The graph shown in the figure applies only if we take it for granted that John is a normal male adult. If it is possible that John is something else, say a child or a horse, then the payoff will no longer be simply a function of John's height but rather of his height relative to the norm for beings of his type.

Notice that if the assertion "*John is tall*" is made *emphatically* then both the satisfaction if it turns out that he is tall and the disappointment if he isn't will be increased. Thus the payoff function for the assertion made with emphasis is like the original payoff function but "scaled up": *i.e.*, with all ordinates multiplied by some factor greater than 1. Thus the representation of meaning by a payoff function is able to take account of emphasis. Insofar as a change in emphasis may be regarded as producing a different assertion of the same sentence, this means that what we are developing here is really a logic of assertions rather than of sentences.

As we have seen, for any given assertion the payoff function is the same for every agent and represents the meaning (determined by society) of the assertion. In this respect the payoff function behaves like the truth function of a sentence in classical logic. To see that this is more than an analogy let us define a *crisp assertion* to be an assertion a whose payoff function takes at most two values; and call it (λ, μ)-*normalized* if these values are λ and μ ($\lambda < \mu$). Then to every (λ, μ)-normalized crisp assertion a there corresponds a unique classical sentence $\theta(a)$ which is true in the world states ω where $a(\omega) = \mu$ and false in those where $a(\omega) = \lambda$. Moreover, for each pair $\{\lambda, \mu\}$ with $\lambda < \mu$ the mapping θ defines a one to one correspondence from the (λ, μ)-normalized crisp assertions onto the sentences of classical logic.

The above considerations suggest that we can pass from classical logic to a *logic of assertions* by making the replacements:

sentence $\longrightarrow$ *assertion* (an act)

truth value $\longrightarrow$ *payoff value* (a utility)

In the following we'll develop a logic based on these replacements. Let's first consider the effect on *connectives*.

3 Connectives in The Logic of Assertions

In propositional classical logic we have the familiar conectives $\neg$ (not), $\wedge$ (and), $\vee$ (or), $\rightarrow$ (implies), $\leftrightarrow$ (is equivalent to). The first is a *unary* connective, a mapping from $\mathcal{S}$ into $\mathcal{S}$, where $\mathcal{S}$ is the set of all sentences; the others are *binary*, mappings from $\mathcal{S}^{\epsilon}$ to $\mathcal{S}$. In addition, the constant sentences T (always true) and F (always false) may be reckoned as 0-ary connectives. Each n-ary connective $\star$ is associated with an n-ary function $\tilde{\star}$ in the set of truth values $\{0, 1\}$, *i.e.*, a function on $\{0, 1\}^n$ to $\{0, 1\}$, the truth value of a composite sentence formed with this connective being obtained from those of its components by applying the function $\tilde{\star}$ to the truth values of the components. In other words, the truth *function* of the composite sentence is obtained from the truth functions of its components by applying the function $\tilde{\star}$ "pointwise on Ω ". For example, for any sentences a and b,

$$(\neg a)(\omega) = \tilde{\neg}(a(\omega)),$$

$$(a \wedge b)(\omega) = (a(\omega))\tilde{\wedge}(b(\omega)),$$

for every world state ω. The functions $\tilde{\neg}$, $\tilde{\wedge}$, ... are given, usually via "truth tables", by the equations,

$$\tilde{\neg}(v) = 1 - v,$$

$$v \tilde{\wedge} w = \begin{cases} 1 & if\ v = w = 1 \\ 0 & otherwise \end{cases}$$

and so on, where v and w denote arbitrary truth values in $\{0, 1\}$.

Mathematically, each connective $\tilde{\star}$ is defined by giving the corresponding function $\tilde{\star}$. So there are four unary connectives, sixteen binary connectives, 256 ternary connectives, and so on; but they can all be expressed in terms of a chosen few "basic connectives". The choice of basic connectives can be made in many ways. Usually, some of the connectives, $\neg$ (NOT), $\wedge$ (AND), $\vee$ (OR), $\rightarrow$ (IMPLIES), $\leftrightarrow$ (IS EQUIVALENT TO), are chosen, not for any special mathematical reasons but just because simple sentences formed with these connectives correspond approximately to certain common natural language statements. This correspondence — which is rather poor and should never be regarded as *defining* the connectives — is indicated by the familiar association of the connectives with the words, "not", "and", "or", *etc.*

Now let us consider the possibilities for connectives in the logic of assertions. Since the analogue of a truth value in the logic of assertions is a utility value, *i.e.*, a real number, one might imagine at first sight that there would be an n-ary connective corresponding to every function on $\Re^n$ to $\Re$[3] However, like length or mass, measurements of utility gain are arbitrary up to a scale factor (choice of unit of utility) and this restricts the functions that can be admitted. Suppose, for example, that $\star$ is a binary connective arising from a binary function $\tilde{\star}$ in $\Re$. For any assertions a and b the assertion $a\star b$ should, of course, be completely determined by the assertions a and b. Certainly it should not be altered by a change in the unit of utility. Now, if (for some world state ω) the payoffs for a and b are u and v then $a\star b$ has the payoff $w = u\tilde{\star}v$. But if we carried out the calculation using a different utility scale in which all utilities are multiplied by λ ($\lambda > 0$) then we would compute *for* $a\star b$ the utility $(\lambda u)\tilde{\star}(\lambda v)$. Since this must be equal to λw, the function $\tilde{\star}$ must satisfy:

$$(1) \quad (\lambda u)\tilde{\star}(\lambda v) = \lambda(u\tilde{\star}v),$$

for all numbers u and v and all $\lambda > 0$. Similarly, we can show that if $\star$ is any unary connective arising from a unary function $\tilde{\star}$ in $\Re$ then

$$(2) \quad \tilde{\star}(\lambda u) = \lambda\tilde{\star}(u),$$

for all u and all $\lambda > 0$. Similarly for 0-ary connectives, *i.e.*, constant-valued assertions: in classical logic there are two, T and F; in the logic of assertions any admissible constant must be invariant under change of utility scale, so there is only one: the function $\mathbf{0}$ given by $\mathbf{0}(\omega) = 0$ for every state ω.

We shall call a function $\tilde{\star}$ with the properties (1) or (2) *admissible.*[4] Some functions in $\Re$ are clearly inadmissible: for instance, the unary function *square*, given by $\tilde{\star}(u) = u^2$, and the binary function, *product*, given by $u\tilde{\star}v = uv$. But many simple unary and binary functions in $\Re$ are admissible. Some examples are:

Scaling:	$\tilde{\star}(u) = ku \ (k > 0)$,
Negative:	$\tilde{\star}(u) = -u$,
Plus:	$u\tilde{\star}v = u+v$,
Minus:	$u\tilde{\star}v = u-v$,
Maximum:	$u\tilde{\star}v = u \vee v = max(u, v)$,
Minimum:	$u\tilde{\star}v = u\wedge v = min(u, v)$.

[3] $\Re$ denotes the set of all real numbers.

[4] In elementary algebra such a function is said to be "positive homogeneous of degree 1".

We shall call these six[5] functions, together with the "0-ary function" $\mathcal{O}$, *basic functions*, and the connectives to which they give rise *basic connectives.* Of the basic connectives the first five are the *linear connectives* and the last two, maximum and minimum, the *lattice connectives.* Of course, the basic functions are not all independent. For instance, *minus* and *minimum* can be expressed in terms of the other four basic functions: $u-v = u+(-v)$, $u \wedge v = -(-u \vee -v)$.

What other admissible functions are there? Well, any function that can be expressed (by means of composition) in terms of basic functions is also admissible. For instance *modulus*, $|u| = u \vee (-u)$, is admissible, and so is any linear combination $ku + k'v$ (where k and k' are any real numbers). This raises the question, Can every admissible function be expressed in terms of basic functions? This is essentially the question of the *truth-functional completeness* of the logic.

In the case of classical logic every function on $\{0, 1\}^n$ to $\{0, 1\}$ determines a connective. Thus all these functions (for every n) are admissible, and, as was noted above (see, for instance, [Rescher69; pp. 62–66]), every connective can be expressed in terms of (suitably chosen) basic connectives. So the answer to the question in the classical case is "yes".

It turns out that in the logic of assertions the answer is again essentially "yes". To express this precisely take the case $n = 2$. One can show, by an application of the Stone-Weierstrass theorem [Dunford&Schwartz58], that if $\tilde{\star}$: $\Re^2 \to \Re$ is any binary admissible function then, given any positive number ϵ, no matter how small, and any number N no matter how large, there is a function $\tilde{\circ}$: $\Re \to \Re$, built by composition from basic functions, that agrees with $\tilde{\star}$ to within ϵ at all points distant no more than N from the origin in $\Re^2$: *i.e.*, $|u\tilde{\star}v - u\tilde{\circ}v| < \epsilon$ whenever $\sqrt{u^2 + v^2} < N$. Since the number of admissible functions is uncountable but only a countable number of functions can be built from the basic functions, and since the discrepancy $|u\tilde{\star}v - u\tilde{\circ}v|$ is positive homogeneous as a function of u and v, it is clear that this is the strongest result that could be hoped for.

4 Classification of States of Belief

As we have seen, there are an infinite number of connectives in the logic of assertions, each being represented by some, generally unary or binary, real-valued function. Among them the ones we have distinguished as "basic"

[5]For simplicity, we shall regard the function, *scaling*, as a single function, although it is really a different function for each value of k.

are singled out only by the simplicity of their functional representation. Although this appears at first to be a purely mathematical property, the fact that the truth values of assertions are utility values, and so play a vital role in determining the conditions under which the assertions would be made, indicates that mathematical simplicity is likely to correspond to simplicity of meaning. This suggests that one should examine the interpretation of simple formal assertions in order to understand the meanings they carry.

Before we can proceed with this task, however, we must consider how an agent, given a payoff function, decides whether to make the corresponding assertion. We assume, of course, that the agent is "rational" in some sense. In particular, his decision whether to make an assertion or not is determined by the interpretation of the values of the payoff function as utilities.

One extreme case (case of perfect information) is easily dealt with: suppose the agent thinks he is omniscient in that he (thinks he) knows, the exact value ω of the present world state. Then, since utility is a measure of preference, he will certainly make the assertion a if $a(\omega) > 0$, and he will certainly not make it if $a(\omega) < 0$. We shall refer to such an agent as a *confident classical agent*, and to ω *as the corresponding world state*. The term "confident" is necessary since we shall introduce *uncertain* classical agents below; and the term "classical" is used since in the domain of crisp assertions, the beliefs of this type of agent can be directly represented by means of classical logic. Indeed, classical logic operates under the presumption that every sentence is either true or false, which corresponds in this domain with the beliefs of a confident classical agent.

It is clear that in real life no agent is, or even thinks he is, omniscient. Consider, then, an agent who recognizes that he has only an imperfect idea of what the world state is. To him, provided he is rational, but regardless of his beliefs, an assertion a will certainly be *acceptable*, in the sense that he is willing to make the assertion, if it happens that $a(\omega) > 0$ for all ω; and it will certainly not be acceptable if $a(\omega) < 0$ for all ω. Whether it is acceptable or not in other cases depends on the agent's *state of belief*: *i.e.*, on exactly what information he (thinks he) has concerning the actual state of the world.

At this point we have to take some position regarding the kinds of "uncertain belief" we are prepared to consider as rational. Two kinds are already relatively familiar: the first, and most well known, is met in the Bayesian approach to decision theory. There the state of belief of any rational agent is represented by an exactly specified (subjective) probability distribution μ over the set Ω of all world states. According to the accepted theory, the agent's choice between available acts is determined by their expected payoffs with respect to this probability distribution. In particular, he will make an assertion a iff (= if and only if) the expected payoff $\int_\Omega a d\mu$, which we shall with a slight abuse of notation denote $\mu(a)$, is ≥ 0. We will refer to an agent whose behaviour can be explained in this way as a confident Bayesian agent. The term "confident" is in contrast to "uncertain", introduced below; the term "Bayesian" reflects the representation of belief in terms of subjective probability.

The second kind of uncertain belief is unrelated to probability. Consider an agent who believes that the actual world state is a member of a certain subset Δ of the set Ω of all world states, but has no idea which state in Δ it is. Then if (as we will assume) the agent takes a conservative view, he will consider an assertion a acceptable iff $a(\omega) \geq 0$ for all ω in Δ, for only in this case will he be sure of not losing in making the assertion. We shall see below that — at least in the case of crisp assertions — a belief of this sort can still be handled by means of classical logic. For this reason we shall refer to an agent with this kind of belief as an *uncertain classical agent.* Such an agent behaves as though he has a team of advisors who are classical agents, one corresponding to each point in Δ, and he makes an assertion if and only if it is approved by every one of his advisors.

A confident classical agent is a particular case of a confident Bayesian agent, the case where the probability is a *point probability distribution*: *i.e.*, one which is concentrated on a single point. A confident classical agent is also a special case of an uncertain classical agent, that in which the set Δ is a singleton. Thus confident Bayesian agents and uncertain classical agents are generalizations in different directions of confident classical agents. It is natural to seek a common generalization of which these are particular cases. This is given by the concept of an uncertain Bayesian agent: By an *uncertain Bayesian agent* we mean an agent who conforms to the Bayesian philosophy, but is unable to decide exactly what probability distribution over Ω he should use to represent his beliefs. Any such agent A may be represented by a set of "possible probability distributions", *i.e.*, by a subset K_A of the set Σ of all probability distributions over Ω. As in the classical case, we shall assume that an uncertain Bayesian agent A behaves conservatively: he will make an assertion a only if the expected payoff is nonnegative for every probability distribution in K_A. An uncertain Bayesian agent thus behaves as though he had a team of confident Bayesian advisors: he makes an assertion if and only if all his advisors recommend it. This picture is very convenient as a way of establishing the properties of an arbitrary uncertain

Bayesian agent. Of course, a confident Bayesian agent is a special case of an uncertain Bayesian agent, that in which the set K is a singleton. Similarly, an uncertain classical agent is a special case of a uncertain Bayesian agent, that in which K is composed entirely of point probability distributions.

We shall not consider any states of belief more general than that exhibited by an uncertain Bayesian agent. This restriction may seem at present to be somewhat *ad hoc*. However, as we shall see in §6, there are good reasons for it. In fact, it can be shown to follow from rather weak *axioms of rationality*, quite independent of the notion of probability, that every rational agent is an uncertain Bayesian agent. In the generalized utility theory developed in [Giles] still weaker axioms of rationality are imposed, with the result that richer states of belief can be discussed without their being considered irrational. These generalizations lead to a more powerful language of assertions. However, for the present we will for simplicity employ only the standard theory.

In establishing results in the following, we therefore assume an arbitrary agent to be an uncertain Bayesian, and we normally make use of the picture of such an agent as reacting to the opinions of a team of confident Bayesian advisors. In this way our conclusions can be deduced from the relatively familiar behaviour of a confident Bayesian agent.

5 Interpreting The Connectives

We are now in a position to discuss the meaning of the connectives, *i.e.*, the way in which the conditions of assertability of a compound assertion are related to those of its components. Note that it is not a question of *assigning* meanings to these assertions: the meanings are already there, determined operationally via decision theory by the interpretation of the payoff values as utilities. This interpretation determines how an assertion employing the connectives is used (by a rational agent) in any given circumstances. The meaning in natural language terms of the assertion is then revealed if we can discover a natural language statement which would be used in the same way.

For the discussion, let Ω be an arbitrary set of *world states*, and let L be a language of assertions on Ω, by which we mean a set of real-valued functions on Ω, closed under the action of the basic connectives. The elements of L are *assertions*, or the *payoff functions* of assertions — for the purposes of this section it is convenient to identify an assertion with its payoff function. In the mathematical theory it is assumed for simplicity that each element of L is a *bounded* function on Ω (but it is not assumed, of course, that every bounded function belongs to L). A topology is assigned to Ω and certain "ideal points" are adjoined, to give a compact space $\tilde{\Omega}$ on which the elements of L are represented as continuous functions. The "probability measures on Ω " referred to in the present account are really probability measures on the Baire sets of $\tilde{\Omega}$. To improve the readability of the present account these and other mathematical details have been suppressed.

Let us first consider the unary connective $-$ (negative). Clearly for a confident Bayesian agent the expected payoff for $-a$, where a is any assertion, is given by $\mu(-a) = -\mu(a)$. Consequently the agent will always be willing to assert either a or $-a$, but never both except in the special case when $\mu(a) = \mu(-a) = 0$ meaning that both a and $-a$ are marginally assertable. This suggests that the connective, negative, corresponds roughly to the negation of classical logic and common language. This impression is supported by the fact that $-(-a) = a$ This equation simply means that the two sides have the same payoff function. We will see more evidence supporting the correspondence between negative and negation below. But see [Giles88] for an example showing that common language negation does not always correspond to the connective, $-$. Of course, in the case of an uncertain Bayesian agent it may happen that, for an assertion a, neither a nor $-a$ is acceptable. This is nothing new: indeed, the same applies in the case of an uncertain agent in classical logic (see §6).

Next consider the unary connective k, where k is an arbitrary positive number. We have already seen that multiplication by k changes only the *emphasis* associated with an assertion. Now, if a is any assertion then ka is acceptable to a confident Bayesian agent if and only if a is acceptable, for $\mu(ka) = k\mu(a)$ for every probability measure μ. But it should not be inferred that ka and a are equivalent insofar as their interpretation is concerned. Indeed, the payoff for ka is k times the payoff for a, so that (if $k > 1$) ka is in a sense "more acceptable" than a. To give this statement more substance let us call $\mu(a)$ the (degree of) *acceptability* to the agent of the assertion a, and denote it also $\alpha(a)$. Unlike the situation in classical logic, where a sentence is simply true or false, in the logic of assertions an assertion has, to a confident Bayesian agent, a definite *degree* of acceptability; it is *acceptable* if this degree is ≥ 0 and *unacceptable* otherwise.

Although this notion of acceptability is introduced at first only for confident Bayesian agents, it can be extended also to an uncertain Bayesian agent A by defining the acceptability to A of an assertion to be the minimum (more strictly the *infimum*) of the acceptabilities assigned to it by his confident Bayesian advisors. An assertion is then acceptable to A precisely iff its acceptability is ≥ 0. Note that this accords with our understanding that an assertion is acceptable to A if and only if it is acceptable to all his advisors.

The meaning of the concept of acceptability is made clearer by considering the case of a crisp assertion a that takes payoff values 0 and 1 (see the end of §2). For a confident Bayesian agent the acceptability of a is just the probability (for him) that the corresponding classical sentence $\tilde{a} = \theta(a)$ is true. For an uncertain Bayesian agent there are two probabilities associated with $\tilde{a}$: an *upper probability* p^+ and a *lower probability* p^-. These are given in terms of acceptabilities by $p^- = \alpha(a)$, $p^+ = -\alpha(-a)$.

Since a (confident or uncertain) classical agent is a particular case of an uncertain Bayesian agent we have now defined the notion of acceptability for all types of agent. Note that the acceptability of an assertion a to a confident classical agent is just the payoff $a(\omega)$, where ω is the corresponding world state. In the case of a crisp assertion that takes payoff values 0 and 1, it coincides with the classical truth value of the corresponding sentence.

Let us next discuss the binary linear connectives, $+$ (plus) and $-$ (minus). Since for a confident Bayesian agent the acceptability $\alpha(a)$ of an assertion a is given by a probability distribution μ, $\alpha(a) = \mu(a)$, acceptability is a linear function on the set L of all assertions: *i.e.*,

$$(3) \qquad \alpha(a+b) = \alpha(a)+\alpha(b) \text{ and } \alpha(ka) = k\alpha(a),$$

for any assertions a and b and for every real number k. For an uncertain Bayesian we can then deduce, using the fact that the acceptability is the infimum of the acceptabilities assigned by his confident Bayesian advisors, that

$$(4) \qquad \alpha(a+b) \geq \alpha(a)+\alpha(b) \text{ and } \alpha(ka) = k\alpha(a),$$

for any assertions a and b and for any *nonnegative* number k. That the statements in (4) cannot be strengthened to the forms given in (3) is shown by a simple example. Suppose A is an uncertain Bayesian agent with two confident Bayesian advisors, and suppose there is an assertion a such that one advisor finds a unacceptable while the other finds $-a$ unacceptable. Then, for A, both $\alpha(a)$ and $\alpha(-a)$ are negative, while, of course, $\alpha(a+(-a)) = \alpha(\mathbf{0}) = 0$.

We can apply this result to the interpretation of the connective, $-$ (minus). Replacing b by $b - a$ in (4) we obtain

$$(5) \quad \alpha(b-a) \le \alpha(b) - \alpha(a).$$

It follows that, for an arbitrary agent for whom $b-a$ is acceptable, b will be acceptable if a is acceptable. So if an agent makes the assertion $b-a$ then we can conclude that if he finds a acceptable then he would find b acceptable too. Thus $b-a$ may be described in the language of classical logic as an assertion that "a implies b". But it is more than this. First, if $b-a$ is acceptable then, by (5), $\alpha(b) \ge \alpha(a)$: *i.e.*, b is *at least as acceptable as* a. Thus if $b-a$ is acceptable then the stronger the acceptability of the premise a the stronger is that of the conclusion b. It is clear that this result (which naturally cannot be represented in classical logic) represents the usual situation in practical life. Observe that here we cannot conclude that the acceptability of the conclusion is *equal* to that of the premise. This is natural, since its acceptability might be increased by other evidence not mentioned in the premise.

It may sometimes be the case that the acceptability of a conclusion b is increased only slightly by an increase in acceptability of the premise a. This type of implication is easily represented in the present formalism by the assertion $b-ka$, where $0 < k < 1$. For we have $\alpha(b-ka) \le \alpha(b) - k\alpha(a)$ which shows that if $b-ka$ is acceptable then we can only conclude that the acceptability of b is at least k times the acceptability of a.

To illustrate further the properties of the new implication suppose that $b-a$ is unacceptable, but only *slightly* unacceptable — say $\alpha(b-a) = -\epsilon$, where ϵ is small and positive. In this case we can conclude that $\alpha(b) \ge \alpha(a) - \epsilon$: *i.e.*, the acceptability of the conclusion can only be slightly less than that of the premise. Again, there is no way in which this natural conclusion can arise as a property of implication in classical logic.

As a further illustration of the logic and to introduce the connective, $+$ (plus), consider the natural language assertion $\mathcal{A}$ = "*If Joan has good grades and good references then she will get the job*". Let a, b, c be respectively the assertions "*Joan has good grades*", "*Joan has good references*", and "*Joan will get the job*", and consider the assertion $f = c - a - b$. Observe that it doesn't matter whether we write f as $(c-a)-b$ or $c-(a+b)$ (or several other algebraically equivalent forms); since the connectives operate on the payoff functions pointwise in Ω, algebraically equivalent expressions are logically equivalent. Taking the second of these forms, we see that f may be read briefly as "$a+b$ implies c". Note that, by (5) and (4),

$$(6) \quad \alpha(f) \le \alpha(c) - \alpha(a+b) \le \alpha(c) - \alpha(a) - \alpha(b).$$

It follows that if f is acceptable then

$$(7) \quad \alpha(c) \geq \alpha(a+b) \geq \alpha(a)+\alpha(b):$$

i.e., the acceptability of c is at least as great as the sum of the acceptabilities of a and b. As an immediate result this means that if a and b are both acceptable then so is c, which indicates that f may be taken as a statement of the form "(a and b) implies c". This shows first that f may tentatively be considered as a formal representative of the natural language assertion $\mathcal{A}$, and secondly, that — in this context at least — $a+b$ serves as a kind of conjunction. As in the case of the connective, minus, this "conjunction" has a number of properties distinguishing it from the conjunction of classical logic. One of these is evident already from (6): insofar as the acceptability of c is concerned a decrease in acceptability of a can be compensated by an increase in acceptability of b: *i.e.*, Joan can compensate for inferior grades by superior references. Now, something of this nature is implicit in most common language statements that have the form of $\mathcal{A}$. Insofar as this is the case, the use of plus may be justified as a way of representing natural language conjunction in the premise of an implication. A second property suggesting that the connective, plus, is a form of conjunction is also evident from (6): if a and b are both acceptable then so is $a+b$. On the other hand, unlike the situation with classical conjunction, it is not the case that if $a+b$ is acceptable then so are both a and b: it is even possible (see the example following (4)) that neither a nor b is acceptable.

Lastly, let us consider the lattice connectives $\vee$ (maximum) and $\wedge$ (minimum). In $\wedge$ we have a much stronger form of conjunction. Since, for any assertions a and b, $a \wedge b \leq a$ and $a \wedge b \leq b$, it follows that for *any* agent if $a \wedge b$ is acceptable then so are a and b. Conversely, in the case of a confident classical agent we clearly have $\alpha(a \wedge b) = min(\alpha(a), \alpha(b))$ which shows that if a and b are both acceptable then so is $a \wedge b$, and it is easy to deduce that the latter holds also in the case of an uncertain classical agent. On the other hand, for a probability measure μ we only have the inequality $\mu(a \wedge b) \leq min(\mu(a), \mu(b))$. As a result, in the case of a confident Bayesian agent a and b may both be acceptable while $a \wedge b$ is not. As an example, suppose a gives payoffs of 2 and –1 respectively according to whether it is raining or not, while the corresponding payoffs for b are –1 and 2 respectively. Then $a \wedge b$ is certain to give a payoff of –1, and so to be unacceptable, but a Bayesian agent who believes the probability of rain to be near 1/2 will find both a and b acceptable. Similar results apply to $\vee$, which acts as a strong form of disjunction.

As one might expect, the lattice connectives $\vee$ (maximum) and $\wedge$ (minimum) are related to the familiar classical connectives denoted by the same symbols, $\vee$ (OR) and $\wedge$ (AND). In fact, it is clear that, given any pair of numbers $\{\lambda, \mu\}$ with $\lambda < \mu$, the mapping θ of the subset $C_{\lambda,\mu}$ of L consisting of all (λ, μ)-normalized crisp assertions onto the set of all sentences of classical logic (see the end of §2) preserves the connectives $\vee$ and $\wedge$: $\theta(a \vee b) = \theta(a) \vee \theta(b)$ and $\theta(a \wedge b) = \theta(a) \wedge \theta(b)$.

Using the mapping θ we can show that the logic of assertions (LOA) is a true extension of classical logic (CL), in that LOA contains CL as a "sublogic". There are many ways of doing this, each starting by fixing a particular scale of utility. Probably the simplest is to take $\lambda = 0$, $\mu = 1$. The set $C_{0,1}$ consists of (the payoff functions of) all the $\{0, 1\}$-valued crisp assertions in L. Each such function can be interpreted as the truth function of a sentence of CL. Now, in $C_{0,1}$ there is a greatest element, the function $\mathbf{1}$ given by $\mathbf{1}(\omega) = 1$ for every world state ω. Let us define in L a connective $\neg$ by the rule: $\neg a$ *is an abbreviation* for $\mathbf{1}-a$, (for every assertion a). It is easy to see that $C_{0,1}$ is closed under the action of the connectives $\wedge$, $\vee$, and $\neg$; and that $C_{0,1}$, equipped with these connectives only, is isomorphic to CL (*i.e.*, to a form of CL in which $\wedge$, $\vee$, $\neg$ are taken as basic connectives). An alternative, which is in some ways nicer, is to take $C_{-1,1}$ instead of $C_{0,1}$ and simply define $\neg$ to coincide with the unary connective, negative. If we associate the payoff values 1 and –1 with the truth values TRUE and FALSE we again get a realization of the same form of CL. In these or many other ways we can see that the logic CL can be obtained by taking part of the set L of all assertions of LOA and a subset of the set of all connectives of LOA.

It may be worth noting at this time that *fuzzy logic*, in its most usual form, infinite-valued Lukasiewicz logic, can also be obtained in a natural way as a sublogic of LOA. Let $F_{0,1}$ denote the subset of L consisting of all assertions whose payoff functions take only values in the closed interval [0, 1]. Define $\neg$ as before ($\neg a$ means $\mathbf{1}-a$) and define a new binary connective $\rightarrow$ by: $a \rightarrow b$ is an abbreviation for $\mathbf{1} \wedge (\mathbf{1}-a+b)$. Then $F_{0,1}$ is closed under the action of the connectives $\wedge$, $\vee$, $\neg$, and $\rightarrow$; and $F_{0,1}$ equipped with these connectives is a realization of the fuzzy logic of all fuzzy sentences (about the world). This shows that, technically at least, LOA is a generalization of fuzzy logic, in that it contains, in a fairly natural way, a realization of fuzzy logic. Whether the interpretation of the family of assertions $F_{0,1}$ agrees with that of the corresponding fuzzy sentences can hardly be decided, since no clear interpretation of the sentences of fuzzy logic has yet been accorded general approval.

In the logic of assertions, the lattice connectives, $\vee$ (maximum) and $\wedge$

(minimum), are more difficult to handle than the linear connectives for the following reason. In the case of a (certain) classical agent the acceptability of any assertion a is the value of $a(\omega)$ for the corresponding world state ω. Since the connectives act pointwise in Ω we have as an immediate consequence:

> (8) *The acceptability of any compound assertion is determined by the acceptabilities of the components.*

In the case of a confident Bayesian agent, the acceptability of an assertion a is the expected payoff $\mu(a)$, where μ is the corresponding probability distribution on Ω. Because μ is a linear function, (8) still holds for the linear connectives, as we saw above. However, it does not hold for the lattice connectives. Suppose, for instance, that a and b are crisp assertions, taking for each world state either the value 1 (TRUE) or 0 (FALSE), and let $\tilde{a}$ and $\tilde{b}$ be the corresponding classical sentences. Then $\mu(a)$ is just the probability that $\tilde{a}$ is true, $\mu(b)$ is the probability that $\tilde{b}$ is true, and $\mu(a \wedge b)$ is the probability that the classical conjunction $\tilde{a}$ AND $\tilde{b}$ is true. But of course this is not determined by the probabilities of $\tilde{a}$ and $\tilde{b}$; it depends also on the correlation between these events. For instance, if $\mu(a) = \mu(b) = 1/2$ then we have $\mu(a \wedge b) = 1/2$ in the case of *maximum correlation* (when $\tilde{a}$ is true when and only when $\tilde{b}$ is true) but $\mu(a \wedge b) = 0$ in the case of minimum correlation (when $\tilde{a}$ is true if and only if $\tilde{b}$ is false) and any intermediate value in other cases. $\mu(a \wedge b) = 1/4$ corresponds to the case of independence of $\tilde{a}$ and $\tilde{b}$. If we regard the acceptability $\alpha(a)$ of an assertion a as a sort of subjective truth value relative to a given agent then we can describe the situation as follows. The logic is truth-functional for all connectives in the case of a confident classical agent, but only for the linear connectives in the case of a confident Bayesian agent. It is not truth-functional even for the linear connectives in the case of an uncertain (classical or Bayesian) agent, but this doesn't matter, as we shall see below. It is for this reason that the lattice connectives are harder to deal with than the linear connectives.

6 Further Developments

In this final section a brief discussion is given of some further developments in the logic of assertions. We start by considering the question of providing satisfactory foundations for this logic.

In introducing the logic of assertions in this paper various assumptions were made in order to simplify the presentation. In particular, in §4 we introduced, under the guise of "agents" of various types, a number of kinds of "uncertain belief". Although in each case the behaviour described could be considered rational, no justification was given for the claim that all reasonable behaviours had been covered. In addition, as presented, the theory appeared to depend on the notions of Bayesian statistics and decision theory, and in particular on probability and utility — concepts which, although plausible, seem inappropriate as a foundation for logic.

We shall now see that the assumptions made are not as arbitrary as they seem. In fact, the structure they lead to follows from certain very simple and primitive axioms, axioms which can (appropriately used) provide a foundation simultaneously for the logic of assertions and for decision theory. In the process of doing this we shall extend to the logic of assertions the concept of *logical consequence*, familiar in classical logic (*CL*). We start by considering, first from the point of view of *CL*, the problem faced by an agent who wishes to describe his state of belief.

In §4 two kinds of "classical agent" were introduced, and in §5 we related these notions to classical logic. A *confident classical agent* is one who knows (or thinks he knows) the exact value ω of the current world state, and hence the truth value (TRUE or FALSE) of every sentence, while an *uncertain classical agent* is one who believes that the actual world state is some unknown member of a known subset Δ of the set Ω of all world states. The state of belief of a confident agent can be represented by the corresponding world state ω. This state determines (and may, for practical purposes, be identified with) the mapping that assigns to each sentence a the truth value $\alpha(\omega)$, where $\alpha(\,\cdot\,)$ is the truth function of the sentence a. Note that this mapping is a *valuation* in that it preserves all the classical connectives. Using this interpretation of a world state, the relation $\models$ of logical consequence in classical logic can be defined: one says a sentence a is a logical consequence of a set of sentences Γ, and writes $\Gamma \models a$, if a is true in every world state in which every sentence in Γ is true; in other words, a is a logical consequence of Γ if a is true in the opinion of every confident classical agent who believes that every sentence in Γ is true.

Classical logic can also represent the beliefs of other kinds of classical agents. In fact, this is the purpose of the machinery of connectives. For instance, by asserting $a \to b$ to be true (a and b being sentences) an agent declares his belief that either b is true or a is false (perhaps both). Such an assertion is useful only in the case of an uncertain agent; there would be no point in making it if the agent already knew the truth values of the component sentences involved. The agent can express other aspects of his belief by making further assertions. In fact, as far as classical logic is concerned, the complete *state of belief* of an uncertain agent may be determined by giving the set of all sentences that he believes to be true. Let us call a set of sentences that arises in this way a *belief set.* Note that the belief set of a confident classical agent is just the set of all sentences that are true in the corresponding world state ω. Thus to say $\Gamma \models a$ is just to say that the sentence a lies in S whenever $S \supseteq \Gamma$ and S is the belief set of a confident classical agent. For convenience. we shall describe a set which is the belief set of a confident classical agent as a "confident classical belief set". Similarly, the terms, "uncertain classical belief set", "confident Bayesian belief set", *etc.*, refer to belief sets of an agent of the corresponding type.

For an idiot, a belief set could be quite arbitrary. However, classical logic considers only "rational" agents, where for classical logic an agent is deemed *rational* if and only if his belief set is closed under the taking of logical consequences: for every sentence a and set of sentences Γ, if $\Gamma \subseteq S$ and $\Gamma \models a$ then $a \in S$; a set S with this property is often called a *theory.* Now, it is easy to show that this condition for rationality is equivalent to the following: an agent is rational if he behaves as though he believes that the actual world state is one of a certain set Δ of world states, but does not know which one; in the sense that he declares a sentence a to be true if and only if it is true in every world state belonging to the set Δ. It follows that the agents that are considered rational in ordinary classical logic are just those that we have previously described as "uncertain classical agents". Moreover, (since each point in Δ corresponds to a confident classical agent) every uncertain classical belief set is an intersection of confident classical belief sets; and conversely, every such intersection is an uncertain classical belief set.

By following the same approach in the logic of assertions (LOA), we can obtain analogues there of the concepts of belief set and logical consequence. In discussing CL we talked about the two types of classical agents; for LOA we must consider instead confident and uncertain Bayesian agents. A confident Bayesian agent (see §4) is represented not by a specific world state ω but rather by a specific probability measure μ on the set Ω of all world states. Just as, in the classical case, the practical function of the world state ω was to assign to each sentence a a truth value $a(\omega)$, so the function of the probability measure μ is to assign to each assertion a a specific *expected payoff* $\mu(a)$; and — just as the truth value of a sentence determines when it can be correctly asserted — so the expected payoff has the property that a is acceptable iff $\mu(a) \geq 0$. Unlike the situation in the classical case, however, the map $a \mapsto \mu(a)$ does not preserve all the connectives but only the linear connectives. Thus the acceptability of a compound assertion is not in general determined by the acceptabilities of its components. In this practical sense *the logic of assertions* is not *truth-functional.* This fact is particularly evident in the case of (0, 1)-normalized crisp assertions, where, as we saw in §5, the expected payoff coincides with the probability of truth.

In spite of this complication of the logic of assertions, we can define the relation $\models$ of logical consequence in the same way as before: an assertion a is a logical consequence of a set of assertions Γ, written $\Gamma \models a$, if a is acceptable in the opinion of every confident Bayesian agent who believes that every assertion in Γ is acceptable. If in the *LOA* we now define the belief set of any agent to be the set of all acceptable assertions, *i.e.*, assertions that he would be willing to make, then this can be expressed as follows:

(9) An assertion a is a logical consequence of Γ iff a belongs to every confident Bayesian belief set that contains Γ.

Now consider an uncertain Bayesian agent as defined in §4. Such an agent behaves as though he has a team of confident Bayesian advisors and regards an assertion as acceptable if and only if it is acceptable to all of them. From this it follows that the belief set of an uncertain Bayesian agent is the intersection of the belief sets of his (confident Bayesian) advisors; and conversely, every intersection of confident Bayesian belief sets is an uncertain Bayesian belief set. Also (9) now gives:

(10) An assertion a is a logical consequence of Γ iff a belongs to every uncertain Bayesian belief set that contains Γ.

Classical logic is truth-functional: the truth value, and hence the "assertability", of any compound sentence is a function of the truth values of its components. Because of this the properties of the logic can be derived directly from the truth tables for the connectives. The logic of assertions is, as we have seen, not truth-functional in this sense. As a result, the rules governing the use of the language can not be so directly obtained from the definitions of the connectives. However, the connections established above show that they may instead be derived from the properties of belief sets. Let us therefore look at the structure of these sets. In this connection remember that an assertion is, for present purposes, identified with its payoff function. A belief set is therefore a set of real-valued functions on the set Ω of all world states.

Consider first the belief set S of an arbitrary confident Bayesian agent characterized by a probability measure μ on Ω. S consists of all assertions a for which the expected payoff $\mu(a)$ is nonnegative. From this we easily obtain the following properties:

(I) S is *upclosed*: if $a \geq b \in S$ then $a \in S$.

(II) S is *additive*: if $a \in S$ and $b \in S$ then $a+b \in S$.

(III) S is *positively homogeneous*: $\mathbf{0} \in S$, and if $a \in S$ and $k \geq 0$ then $ka \in S$.

(IV) S is (topologically) *closed*: if $\{a_n\}$ is a sequence of assertions in S and $a_n(\omega) \to a(\omega)$ uniformly for all ω in Ω then $a \in S$.

(V) S is *definite*: if $a \notin S$ then $-a \in S$.

Now, an arbitrary uncertain Bayesian belief set is an intersection of confident Bayesian belief sets. Using this fact we deduce easily that, because a confident Bayesian belief set has the property (I), so does an uncertain Bayesian belief set. The same applies to properties (II), (III), (IV). However, simple examples show that an uncertain Bayesian belief set need not have the property (V).

Notice that the properties (I) — (V) have very simple meanings: (I) says if a has, for every world state, at least as large a payoff as b and b is acceptable then so is a; (II) says that if a and b are acceptable then the assertion $a+b$ (which, because the payoff is additive, may be interpreted roughly as an assertion of both a and b) is also acceptable; (III) says that if a can be asserted without danger of loss then it can also be asserted *emphatically* without danger. (II) and (III) can be criticized. See below. (IV) is a typical *closure* or *Archimedean* axiom: if each of a sequence of assertions is acceptable then so is the limit of the sequence. (V) is a little more specialized; it says that if a is not acceptable then its "reverse" $-a$ must be. This is plausible for a confident Bayesian agent: if he would lose in asserting a he would equally gain in asserting $-a$. But an uncertain agent might be afraid of losing in making either of the assertions a and $-a$ and so consider neither acceptable.

The plausibility, for a general agent, of properties (I, II, III, IV) suggests the possibility of taking them not as derived results but as axioms. This gives rise to a new formulation of the logic which is in many respects more satisfactory. The reformulation can be done in two ways. In the first way we start as here with the notion, derived from decision theory, of the payoff function of an assertion. As in decision theory, this is a real-valued function on the set Ω of all world states. Then we introduce the notion of a *belief set* as a set satisfying (I) — (IV), these now being taken as axioms and justified by their plausibility with reference to the understood meaning of a belief set. In this way a foundation for the logic is obtained in which the notion of subjective probability is never used. Nevertheless, it can be proved (surprisingly) that (*a*) every belief set is an intersection of "maximal belief sets" and (*b*) every maximal belief set can be represented uniquely as a "confident Bayesian belief set" given by a probability measure on Ω as described above. In this way we obtain the same structure as described here, in a way which is more satisfactory as a foundation for the logic in that it depends on a smaller number of more fundamental concepts.

The second way of reformulating the theory represents a further step in the same direction. An examination of the axiomatic structure of the first formulation reveals a great deal of similarity to modern formulations of utility and decision theory [Fishburn81]. This is not surprising, since the basic problem of the logic of assertions, whether or not to make a given assertion, is just a special case of the decision problem, which of a given set of possible acts to carry out. As a result, it is possible to give a formulation of the foundations of the logic which, by itself containing the necessary foundations, remains independent of any prior account of decision theory. Work in this direction is currently in progress.

This second approach offers the possibility of overcoming a certain difficulty which was briefly noted above. Of the axioms (I) — (IV) the least well justified are (II) and (III). We shall call these the *homogeneous* axioms, in view of the nature of the mathematical structure which they induce, and the logic of assertions described in this paper will accordingly be called the *homogeneous logic of assertions.* The justification offered above for axioms (II) and (III) can be criticized: in practice it is certainly not the case that an agent who is willing to make an assertion is necessarily willing to make it with an arbitrarily large emphasis; the risk involved in doing so may be too great. In the same way, one may be prepared to place a small bet on some event but unwilling to place a large bet. A similar criticism may be applied to (III). By dropping the homogeneous axioms, or more precisely by replacing them by a weaker "convexity" axiom

if a and b are acceptable then so is $(a+b)/2$

a more general form, the *convex logic of assertions*, is obtained. In this alternative logic more general states of belief can be described than those allowed in the homogeneous case. This means that more general beliefs can be conveyed, so that the resulting language is *richer*. Normally, one would expect that as a result the mathematical structure would be more complicated and the processes of reasoning more difficult. It turns out, however, that the theory is still quite manageable, in fact very little is changed. Roughly speaking, where the homogeneous theory is dominated by convex cones (in linear spaces) the convex theory has convex sets and convex functions. The convex theory is outlined in [Giles88]. This paper also contains a demonstration of the power of *linear logic*, by which I mean the logic of assertions *without* the lattice connectives $\wedge$ and $\vee$. In particular, a number of practical examples are given, indicating the richness of linear logic when the convex language is used.

Now, the difficulty, mentioned above, affecting the convex theory arises when one attempts to give an decision-theoretic foundation for it. Conventional decision theory (expected utility theory) corresponds to the *homogeneous* rather than to the *convex* case. As a result, to serve as a foundation for the convex logic of assertions a generalization of standard expected utility theory is required. A suitable generalization has recently been given [Giles], but its application to a complete decision-theoretic treatment of the logic has yet to be worked out.

Other work that remains to be done in the development of these logics includes the extension to a first order language, and the development of a system of deduction. As indicated above, in the latter direction it is relatively easy to describe a system of deduction for *linear logic*. An outline of such a system was given in [Giles85]. Owing to the dominance of convex sets and functions the process of reasoning reduces there essentially to solving a problem in *linear programming*. It is to be hoped that the powerful techniques developed in this field will prove valuable in the practical applications of the logic of assertions.

Acknowlegements. This work was supported by a grant from the National Science and Engineering Research Council of Canada.

7 Bibliography

Dunford, N. and J. Schwartz. *Linear Operators, Part I*, Interscience, 1958.

Fishburn, P. "Subjective expected utility: a review of normative theories," *Theory and Decision 13*, 139–199, 1981.

Giles, R. "A resolution logic for fuzzy reasoning," *Proceedings of the Fifteenth International Symposium on Multiple-Valued Logic 60–67*, 1985.

Giles, R. "A utility-valued logic for decision making," *International Journal of Approximate Reasoning 2*, pp. 113–141, 1988.

Giles, R. "A generalization of the theory of subjective probability and expected utility," to appear in *Synthese.*

James, W. *Pragmatism*, Longmans Green, 1907.

Luce, R. and H. Raiffa. *Games and Decisions*, Wiley, 1957.

Raiffa, H. *Decision Analysis*, Addison-Wesley, 1968.

Rescher, N. *Many-Valued Logic*, McGraw-Hill, 1969.