CRÍTICA, *Revista Hispanoamericana de Filosofía*
Vol. XXVII, No. 81 (diciembre de 1995): 27–54

# THE LOGIC OF UNCERTAINTY*

DOROTHY EDGINGTON
London University

A good question with which to begin a study of the logic of uncertainty is this: suppose you have an argument which you recognize to be valid —you recognize the impossibility that it have true premisses and a false conclusion; you think, but are not sure that its premisses are true. How, if at all, does this constrain what you should think about the conclusion of the argument? In section 1 I show that there is an instructive and plausible answer to this question, provided we treat uncertainty in terms of probability. There are two parts to this assumption: first, that it makes sense to quantify uncertainty; and second, that the quantities obey what I present as *the* fundamental principle of probabilistic structure: that the sum of the probabilities of a set of exclusive and exhaustive possibilities is 1, the value assigned to a certainty. These two claims are defended in sections 2 and 3. In section 4 I show that the fundamental principle is rich in consequences: all the other principles of probabilistic structure which do not require the

* Although intended as the first chapter of a book, it was suggested to me that this material is sufficiently self-contained to appear usefully as an article. It formed part of a course I gave recently at the Instituto de Investigaciones Filosóficas, Universidad Nacional Autónoma de México. I thank my colleagues there for much useful criticism. I thank also the British Academy for the award of a Research Readership which allowed me to proceed with this work.

notion of conditional probability[1] are shown to be derivable from it.

## 1. *Off-Premisses Logic*

A notice appeared on the board: due to Smith's illness, either Dullard or Bright will give this week's logic lecture. On his way to the lecture, John meets Bright hurrying in the opposite direction.

"Not lecturing today?" he asks.

"No —appointment elsewhere— must dash", says Bright. John performs a little inference, decides that he too would be better off elsewhere, and makes for the pub, where he is greeted by a fellow-student.

"I thought you were a logic fan!", says Mary.

"I am —it's logic that brings me here." He explains.

"But how can you be sure that your premisses are true?" asks Mary, who, though not a logic fan, has taken other philosophical studies to heart.

"I can't be *absolutely* sure. Notice boards are not infallible. And perhaps Bright was lying, or was about to remember his promise to lecture. He may even have a twin brother who teaches down the road.... But, in my judgement, my premisses are very likely to be true. The argument is valid; so the conclusion is very likely to be true."

"Wrong!", says Mary. "A valid argument can have very probable premisses, and a conclusion which is very improbable, or even certainly false. Haven't you heard of the Lottery Paradox? We have many premisses —a million, say: Ticket 1 won't win; Ticket 2 won't win; and so on. The conclusion is that none of the tickets will win. Each pre-

[1] This very important notion is not discussed here.

miss is highly probable, the conclusion is —practically— certain to be false.[2]

"This is an extreme example, but it illustrates a general phenomenon. You admit the notice may be wrong, in which case your conclusion would be false. And even if the notice is correct, you admit that your second premiss may be wrong, in which case, also, your conclusion would be false —ruling out as unthinkable a double act. Your conclusion, then, has more chance of being wrong than either premiss. Validity does not preserve degree of certainty. Logic is useless for reasoning from the sort of information typically available to us —almost invariably, uncertain information."

"But *surely* a conclusion validly derived from only two premisses, each close to certain, can't be much less certain", replies John. But his tone of voice, and his 'surely', betray his doubts about how to pursue the matter. Enter Jane, who is writing a thesis on probability.

"Elementary," says Jane, when she hears the problem, "at least if you are prepared to concede that degrees of certainty should, ideally, obey the rules of probability. Call the uncertainty of a proposition one minus its probability. Then we can establish this: *the uncertainty of a conclusion cannot exceed the sum of the uncertainties of the premisses from which it is validly derived.*

"First, if an argument has one premiss and you recognize it to be valid, you cannot consistently think the conclusion is less probable —more uncertain— than the premiss."

[2] Another problem of the same structure which Mary could have mentioned is the Paradox of the Preface: the author announces that she is sure that her book is not entirely free from falsehood —she is bound to have made some mistake somewhere. Of each individual statement, she believes that it is true. Yet she does not believe that they are *all* true.

(She scribbles on the back of an envelope,[3] sees frowns, then draws two concentric circles, sees nods, and proceeds.)

"Second, a many-premiss argument can be reduced to a single-premiss argument by conjoining the premisses. Third —this is the interesting bit, which I'll convince you of shortly— the uncertainty of $A\&B$ cannot exceed the sum of the uncertainties of $A$ and of $B$. Fourth, conjunction and addition being associative, we can generalize,[4] obtaining" (she scribbles again):

$u(C) \leq u(A_1\& \ldots \&A_n) \leq u(A_1) + \cdots + u(A_n)$, when $A_1, \ldots A_n$ entail $C$.

She takes a matchbox from her pocket, and two broken matches.

"I see you don't like algebra, so here is a model of the third part of the proof. The matchbox has length 1, the

[3] To recognize that $A$ entails $C$ is to give zero probability to ($A$ and not $C$): $p(A\&\neg C) = 0$; so $p(A) = p(A\&C)$. $p(C) = p(A\&C) + p(\neg A\&C) \geq p(A)$. (This is explained more fully in section 4. For the moment, think of a one-premiss valid argument like 'It's square, so it has four sides'; convince yourself that it would be absurd to think the premiss more likely than the conclusion; and that the example is typical.)

[4] That is, $A\&(B\&C)$ is the same as $(A\&B)\&C$; and $(x+y)+z = x+(y+z)$; in each case, the brackets are redundant. It follows that if we have the result for two conjuncts, it can be extended upwards for any number of conjuncts. In the case of three conjuncts, (i) $u(A\&B\&C) = u((A\&B)\&C) \leq u(A\&B) + u(C)$. But $u(A\&B) \leq u(A) + u(B)$; so (ii) $u(A\&B) + u(C) \leq u(A) + u(B) + u(C)$. (i) and (ii) yield the desired result: $u(A\&B\&C) \leq u(A) + u(B) + u(C)$.

matches lengths 0.9 and 0.8. They are to be put lengthwise in the box. What is the minimum overlap of the two matches? Well, placed end to end with no overlap, their length totals 1.7. So to get them in a box of length 1, they must overlap by a minimum of 0.7. (The maximum overlap is 0.8, the length of the shorter match.) Now, let the lengths of the matches represent the probabilities of two propositions, $A$ and $B$, the lengths of the box occupied by just $A$, just $B$, both, and neither, represent respectively the probabilities of $A\&\neg B$, $\neg A\&B$, $A\&B$ and $\neg A\&\neg B$. (These four probabilities must sum to 1: that is the substantive assumption of probabilistic structure.) We now have a model of our result: given $p(A) = 0.9$ ($u(A) = 0.1$); $p(B) = 0.8$ ($u(B) = 0.2$), we have $p(A\&B) \geq 0.7$ or $u(A\&B) \leq 0.3 = 0.1 + 0.2 = u(A) + u(B)$.[5]

"Call John's premisses $A$ and $B$. Suppose he thinks each is at least 99% probable. This commits him to thinking that $A\&B$ is at least 98% probable. But $A\&B$ entails his conclusion $C$ —it's impossible that $(A\&B)\&\neg C$. So, given the values he assigns to his premisses, $C$ must also be at least 98% probable. A valid argument with a few highly probable premisses *does* guarantee a high probability

[5] The importance of this feature of valid arguments was first recognized by Ernest Adams. See his 'Probability and the Logic of Conditionals' in J. Hintikka and P. Suppes (eds.), *Aspects of Inductive Logic*, (Amsterdam: North Holland) 1966; and his *The Logic of Conditionals* (Dordrecht: Reidel), 1975, chapter 2.

for the conclusion —though not quite as high as the premisses. Of course, if there are lots of premisses, or a few but they're only *fairly* probable, the 'conjunctivitis' gets serious: the conclusion can inherit the uncertainty of each premiss. That's just common sense. Still, a valid argument, unlike an invalid one, provides a constraint on the fall in probability; so logic *can* be usefully applied to uncertain information."[6]

John beams, his faith in the power of reason restored. But a sceptical frame of mind, that healthy antidote to faith, soon finds a new locus of doubt:

"But why should I accept that degrees of confidence should behave as probabilities?" Mary asks. "I don't dispute that probability theory has many useful applications —to random sampling, for instance; and to phenomena like coin-tossing which appear to be random but show stable long-run proportions. But why should I believe that degrees of confidence in propositions in general —that Labour will win the next election, that the world began with a Big Bang, that this painting was painted by Goya, and so forth, should obey the laws of probability?"

[6] This principle applies to any argument in which it is impossible that the premisses are true and the conclusion false (and not merely to *formally* valid arguments). Suppose you think it's about 90% likely that Mary will get a higher mark than John in the test; and about 80% likely that John will get a higher mark than Sue. You must then think it is at least 70% likely that Mary will get a higher mark than Sue. On the other hand, suppose, in a "round robin" tennis competition, in which everyone plays everyone, you think it is 90% likely that Mary will beat John, and 80% likely that John will beat Sue; that leaves you free to think what you like about whether Mary will beat Sue —if you think that is 0% likely, there is no inconsistency in your beliefs. (For it is not impossible that Mary beats John, John beats Sue, and Sue beats Mary.)

## 2. *Quantifying Uncertainty*

2.1. A person can be more or less close to certain that a proposition is true. We may express this by saying: there are degrees of belief (or confidence, or certainty). The fact is of some importance about our capacities for recognizing truth. Perception can deliver less-than-certain judgements; so can memory, testimony, and inference, from effect to likely cause, from cause to likely effect. A threefold classification of a person's belief-like attitude towards a proposition —belief, disbelief, neither, ignoring degrees of certainty— is seriously inadequate. Inadequate for what? Inadequate to the main purpose served by having beliefs, and ascribing them to others —that of helping explain and justify actions, and guiding deliberations about what to do. The threefold classification —believes that *p*, disbelieves that *p*, neither believes nor disbelieves *p* (roughly —yes, no, don't know) has two drawbacks. The minor drawback is that the border between the first and third, or the second and third, category is very unclear. Are one's 'beliefs' (1) the things one is certain of, (2) the things one takes as at least close to certain, or (3) the things one judges more likely than not? How much doubt, if any, is compatible with believing (or disbelieving) something? Ordinary usage does not give a clear answer: context, and the importance of the issue, will play a large part in determining which propositions a person is willing to endorse, unqualified, as something she believes. The major difficulty is that, however we settle the minor one, at least one of the categories will be far too crude to explain or justify differences in behaviour, or to guide deliberation about what to do. It will ignore important differences; and also, it will ignore important similarities, between cases which fall on opposite sides of these lines. Take option (3): my beliefs are the things I judge more likely than not. *A* and *B* both believe that there will not be a storm at sea; keep all other relevant factors constant; yet $A$ takes out his boat, $B$ does not take out his boat. Why? Because $A$ is virtually certain that there will not be a storm, while $B$ thinks it just more likely than not. The difference does not come out in terms of what they believe, on this account. $C$ thinks it just less likely than not. $C$ does not believe that there will be no storm; yet there is much less relevant difference between $B$ and $C$ (a believer and a non-believer) than between $A$ and $B$ (two believers). In assimilating importantly different cases, and discriminating between unimportantly different cases, this way of drawing a line ill serves our purposes.

With options (1) and (2), on the other hand, the 'neither' category catches far too much. I need to buy something in a hurry —stamps, say. Speed is the only relevant consideration. I have a choice of two equidistant shops in opposite directions. More often than not, one has to wait longer to be served in shop $A$. So I go to shop $B$. I do not reach the standard of *believing* that shop $B$ will be quicker, or that it will not be quicker (though I am more confident than not that shop $B$ will be quicker). The threefold classification does not explain my choice of shop. Again, someone who believes and someone who falls just short of believing may be more relevantly similar than two who neither believe nor disbelieve. Again, these lines do not serve us well.

One often reads, in works on practical reason, something along the following lines:

> $X$ has a motivating reason to $\phi$ if and only if he desires that $p$ and believes that $\phi$-ing will bring it about that $p$.[7]

This is intended to be compatible with the thought $X$ may have a much stronger motivating reason not to $\phi$. It has the,

[7] See for example Michael Smith, *The Moral Problem* (Blackwell, 1994), pp. 92–93.

to my mind, unfortunate consequence that if I want money and believe that I would get away with stealing some, I thereby have a motivating reason to steal. But the objection in the other direction is more decisive. Suppose there is one known treatment for my disease, which is estimated to have a 40% success rate (and no ill effects). I do not believe that taking the treatment will cure me. But, *contra* the above thesis, I do have a motivating reason to take it. Again, we ignore degrees of certainty at our peril.

I have focused on the need for degrees of belief in practical reasoning. The same could be said about theoretical reasoning, for instance, concerning the evidence for a scientific conjecture. Different people have different evidence, or assess the available evidence differently. Any attempt to sort people into three classes, the believers, the disbelievers, and the agnostics, will classify in an unhelpful manner: two people with almost identical opinions may fall on different sides of these lines; and two people with radically different opinions may fall within one class. The important question of how uncertain you are, is overlooked.

2.2. Defenders of the threefold classification, opponents of degrees of belief, might reply that in the intermediate cases in which I neither believe nor disbelieve, I may nevertheless believe (or disbelieve) something else, namely that *it is probable that* (e.g.) going to shop *B* will be quicker, and this belief (or disbelief), will differentially explain my action. Rather than extend the range of epistemic attitudes to a proposition, they extend the range of propositions believed. Let us call such a person an ‘absolutist’ about belief.

The absolutist owes us an account of the content of these more complex beliefs, one which makes perspicuous their role in guiding behaviour. There is, I allow, more than one interpretation of probability. The degree theorist claims that one legitimate interpretation of a probability judgement is as an expression of a degree of belief in a proposition, and it is this interpretation which is relevant to the explanation of behaviour. In denying this, the absolutist must find another interpretation of 'Probably *p*', which will also play a role in explaining the behaviour of someone with beliefs of this form. It won't do, for example, to interpret 'probably, going to shop *B* will be quicker' as 'usually going to shop *B* is quicker'; for it is possible to believe that while (fully) believing, or disbelieving, that on this occasion shop *B* will be quicker.

Let us allow the absolutist appeal to the notion of objective probability-chance, as a real feature of the world: one may believe that the "real" probability of, say, heads on the next toss of this coin, is, say, 0.5; and, let us concede, this belief will explain and justify the same behaviour as the degree-theorist's supposed belief to degree 0.5 that the coin will land heads.

But this phenomenon is insufficiently general for the absolutist's needs. With the possible exception of a narrow class of "trivially analytic" propositions which cannot be understood without being accepted, *any* proposition can be the subject of uncertainty; but only some concern the possible outcomes of chance processes. Even amongst those which do, some concern the past outcomes of chance processes —the possible outcomes have now come about, or failed to come about (the coin has landed); and our judgements about the truth-value of such propositions need not proceed via knowledge of the chance they *had* of being true. We may have more direct (yet inconclusive) evidence as to how the coin landed.

The ubiquity of uncertainty generates the following crucial difficulty for the absolutist: whatever interpretation he gives of propositional contents of the type 'It is probable that *p*', which may be simply believed, *that propositional content, like any other, can itself be the object of uncer-*

*tainty*. Real chances, for instance, are the sort of things one can be, and typically is, uncertain about. Is this state to be interpreted as belief in a proposition which ascribes a "real" probability to the proposition "It is probable that $p$"? If this can be made sense of, the objection merely iterates: this propositional content, like any other, can itself be uncertain. The prospects of eliminating degrees of certainty are dim.[8]

In addition, the Lottery Paradox and the Paradox of the Preface are serious obstacles to the absolutist conception of belief. For the absolutist, this principle seems compelling: the belief that $p$ and the belief that $q$ may legitimately generate the belief that $p \,\&\, q$. But iteration of this principle many times leads to your believing the conjunction of all statements of the form "Ticket $n$ won't win"; and the conjunction of all the statements in the book. The above principle is harmless if the absolutist interprets belief as entailing complete certainty. So interpreted, almost all propositions are such that we neither believe nor disbelieve them. But if belief does not entail certainty, it is only the degree theorist who can explain why a single application of the above principle is approximately true, but repeated applications allow the conjunction to become less and less certain relative to the beliefs in the conjuncts: repeated applications increase uncertainty.

2.3. Now let us turn to the idea that a person's degree of belief in a proposition (at a time) can be represented by a number. At one extreme, if $X$ is completely sure that $A$, let us say that $b_X(A) = 1$. At the other extreme, $X$ may be completely sure that $A$ is not true; then, let us say, $b_X(A) = 0$. In the middle, $X$ may think $A$ is equally

[8] Let me emphasize that I am not objecting to the notion of objective chance, but merely to the idea that degrees of confidence are eliminable in favour of it.

likely to be true or not true, in which case, let us say, $b_X(A) = 1/2$. Or $X$ may be nearly certain that $A$, think $A$ is a little more likely than not, less likely than not, very unlikely, etc.

Except in particularly amenable special cases, degrees of belief are usually not very precise. But our ordinary judgements do admit of more than mere ordinal comparisons of the form: $A$ is more likely than $B$, which is more likely than $C$. We are capable of judging that $A$ is *much* more likely than $B$; $B$ is *slightly* more probable than $C$, etc. It is important that this is so. Consider two people each facing a problem of the following structure: a course of action may lead to one of three possible outcomes, $A$, $B$ or $C$. Each judges $A$ to be more probable than $B$ and $B$ more probable than $C$. Each judges $A$ to be desirable, $B$ and $C$ to be undesirable (indeed, assuming for the sake of symmetry that this makes sense, each places the same positive or negative values on $A$, $B$ and $C$). But the first judges $A$ to be much more likely than $B$, $B$ much more likely than $C$, while the second has $A$ just a little more likely than $B$, $B$ a little more likely than $C$. It is this difference which may well explain the first's performing the action, the second's refraining from doing so.

For example, take two situations in which there may be three outcomes of deciding to take a certain flight: ($A$) arriving at one's destination on time; ($B$) arriving seriously delayed; ($C$) not arriving at all —a crash. Keep the values and disvalues of the outcomes the same across the two situations, and also the fact that $A$ is judged more likely than $B$, and $B$ more likely than $C$. But suppose that in the first, you judge that $A$ is almost certain; and in the unlikely event that not $A$, $B$ is much more likely than $C$. In the second situation, although the ordering is the same, you judge $A$ only a little more likely than $B$, and $B$ only a little more likely than $C$. It is this latter fact that will explain why you board one plane and not the other.

Thus, our ordinary capacities permit judgements of the form 'the difference in probability between $A$ and $B$ is greater than (less than, roughly equal to) the difference in probability between $B$ and $C$'; we do, implicitly, operate with a rough and ready scale. It is no doubt an idealization to represent a person's degree of belief that $A$ by a number between 0 and 1, the nearer to 1, the closer to certain she is that $A$. Degrees of belief are not usually that precise. But, as I have tried to show, it is not too far removed from our actual practice to prevent its being a useful idealization. We could say that we are modelling human judgement in terms of the judgements of hypothetical beings who are like us except that their degrees of belief are always sharp. And it *is* a useful idealization, giving us access to arithmetic —addition, multiplication, etc.— to exhibit the logical relations between a person's degrees of belief. Of course, caution is required in interpreting the results: interesting results must be robust enough to be independent of small numerical differences, which may be features of the idealization without real significance.

We have already seen an example of the use of numbers in the result of section 1: if an argument is valid, the uncertainty of the conclusion cannot exceed the sum of the uncertainties of the premisses. This yields plausible qualitative results: for a given valid argument with premisses close to certain, the conclusion will be close to certain (though perhaps less close than each premiss individually), provided there are not too many premisses —a large number of small degrees of uncertainty can transmit a large, even a maximal degree of uncertainty to the conclusion. This is just one result. The proof of the pudding will be in the eating, and we have had just one preliminary taste. What matters is whether we can develop a good and useful theory of reasoning with uncertain beliefs in terms of the numerical idealization of uncertainty. The situation is not different in principle from the use of numbers in physics to represent lengths, temperatures, mass, etc. Our measurements of these are always imprecise; but our practice of so representing them provides an indisputably useful model of their behaviour; and this is so independently of whether we believe that the phenomena themselves must have precise numerical values.

## 3. *The Partition Principle*

3.1. What logical principles govern degrees of belief? Logic does not tell you what to believe, but rather that some beliefs rule out others: some combinations of belief are consistent, other combinations are not. The fundamental principle governing degrees of belief, I shall argue, is this: take a set of exclusive and exhaustive possibilities —a set of propositions such that it is impossible that more than one of them is true, and necessary that at least one of them is true. (For simplicity, we restrict ourselves to finite such sets.) Call such a set a *partition*. A person's degrees of belief in the members of a partition must sum to 1, the value to be assigned to a certainty. Let us call this the Partition Principle.

If, in place of 'a person's degrees of belief in', I had said 'the probability of', we have what can serve as the fundamental principle of probability theory. Accepting that degrees of belief obey the Partition Principle amounts to the claim that degrees of belief have the structure of probabilities. All the principles of the logic of probability which do not involve the concept of conditional probability can be shown to follow from this principle.

A partition, in the strict sense, is guaranteed to be such by logic: it is logically impossible that more than one member is true, and logically necessary that at least one member is true. Elementary classical logic yields innumerable examples of partitions, those of the form $\{A, \neg A\}$; $\{A\&B, A\&\neg B, \neg A\}$, for instance.[9] We shall sometimes use, for illustrative purposes, examples which are not partitions in this strict sense —{Heads, Tails} as the possible outcomes of a toss, the set of propositions of the form 'Horse $i$ will win' for the horses, numbered 1 to $n$, in the race. It is not logic which guarantees that the coin won't continue forever upwards, or the race won't end in a dead heat or be called off. These are strictly partitions only relative to a set of background assumptions.

A dilemma arises about what to say when the logic which guarantees that something is a partition is extremely difficult, perhaps yet undiscovered. For example, suppose that $A$ is logically equivalent to $B$, but no one has discovered this abstruse fact. So $\{A, \neg B\}$ is a partition. Are we to be accused of irrationality in having degrees of belief in $A$ and $\neg B$ which do not sum to 1? In a sense yes —we have inconsistent beliefs. But to the extent that we want our principles of rationality to be within our powers, we do not want it to be the case that only the logically omniscient can be rational. I shelve this problem here: all the strict partitions with which we will be concerned are guaranteed to be such by elementary recognizable logical considerations.

For example, the lines of a truth table constitute a partition; and the Partition Principle dictates that one's dis-

[9] Some non-classical logics will deny that these are examples of partitions. Intuitionist logic denies that $A$ and $\neg A$ are always necessarily exhaustive possibilities; and paraconsistent logic, more radically, denies that they are always necessarily exclusive possibilities. Note that neither gives a reason to doubt the Partition Principle, but they dispute about what are examples of partitions. This is not the place to enter into such disputes: for the time being, we shall try to steer clear of the domains in which classical logic is contentious.

tribution of belief (or probability) over the lines of a truth table must sum to 1.

| $A$ | $B$ | (i) | (ii) | (iii) | (iv) | (v) |
|---|---|---|---|---|---|---|
| T | T | 1 | 0.25 | 0.1 | 0.3 | 0.1 |
| T | F | 0 | 0.25 | 0.8 | 0.4 | 0.1 |
| F | T | 0 | 0.25 | 0.1 | 0.2 | 0.1 |
| F | F | 0 | 0.25 | 0 | 0.3 | 0.1 |

In the above table, columns (i), (ii) and (iii) each represent legitimate combinations of degrees of belief in the possibilities represented by the four rows; columns (iv) and (v) represent illegitimate distributions of belief, summing, in the case of (iv) to more than 1, and in the case of (v) to less than 1.

Why accept the Partition Principle? It is hard to find anything more basic in terms of which to argue for it. I shall first try to present it as intuitively compelling; then I shall present an argument which I think has some persuasive force against the sceptic. In the next section, I show that it is powerful in generating important and plausible consequences.

The basic idea behind the principle is this: a set of exhaustive alternative possibilities *compete* for belief. If you give full belief to one member of such a set, you have none left over for any of the others. If you distribute your belief equally over the $n$ members of a partition, you have degree of belief $1/n$ for each of them. The Partition Principle is equivalent to this: for an $n$-membered partition, the average probability of a member must be $1/n$. The principle is silent on how you distribute your probabilities between the $n$ members. But it dictates that if some member gets more than the average, some other member or members get less, to compensate. For example, in a three-horse race (assuming one and only one horse will win), it would be irrational to think that each horse was as likely as not to win.

Suppose someone were to say (as someone once did say to me) in connection with a problem about a horse race:

"Look, this is a really good horse. So it's got a very high chance of winning —and that's a fact which is just about it— which has nothing whatsoever to do with the other horses in the race."

This is a mistake: a horse's chance of winning depends not just on how good it is, but on what the competition is like. And what goes for horses goes for possibilities in general: a strong case for the truth of possibility *A* is undermined by a strong case for the truth of a possibility incompatible with *A*.

3.2. Imagine a sceptic who is unconvinced —Mary, say. "On pain of what do I disobey this principle?", she asks. "In a three-horse race (granting that one and only one horse will win), what is wrong with thinking each horse is as likely as not to win?" Here is an argument designed to answer that question.

We have already said that a primary function of uncertain beliefs is as ingredients in decision-making, and that is where this argument locates them. A simple model of the structure of decision-making under conditions of uncertainty is provided by betting behaviour. The first version of this argument constructs a case where the betting behaviour goes disastrously wrong, and the only plausible explanation of this fact is violation of the Partition Principle.

If someone thinks a proposition is as likely as not, i.e. 50–50, then in normal circumstances it is not irrational for her to accept a bet such that she gains £12 if the proposition turns out to be true, and loses £10 if the proposition turns out to be false. What are normal circumstances? Well, she prefers more money to less: she is not intent on dispersing her embarrassing wealth. And the difference, for her, between being £12 richer and the status quo exceeds the difference between the status quo and being £10 poorer: it would not be a disaster to lose £10; it would not be of negligible value to gain £12.[10] She is not abnormally adversely disposed to risk, or to gambling in particular. Even if circumstances are normal, I need not insist that it is rationally obligatory for someone to accept such a bet —merely that it is rationally permissible that she do so. We would not raise our eyebrows if someone, with such a probability judgement, accepts a bet at these odds: we would not consider the behaviour peculiarly in need of a special explanation.

Now Mary thinks horse *A* is as likely as not to win. She accepts such a bet on the proposition 'horse *A* will win'. She thinks just the same about horses *B* and *C*. Conditions, let us suppose, are still normal. It is rationally permissible for her to accept those bets also. She does so. But now, Mary has ended up giving away £8 whatever happens: whichever horse wins, she gains £12 on one bet and loses £10 on each of the other two.

| | *A* wins | *B* wins | *C* wins |
|---|---|---|---|
| Bet 1 | +12 | −10 | −10 |
| Bet 2 | −10 | +12 | −10 |
| Bet 3 | −10 | −10 | +12 |
| Overall | − 8 | − 8 | − 8 |

[10] One *can* state this argument in terms of "real" values rather than monetary values (and choices of actions other than bets). But it would take us too far afield to justify the structure of "real" values. Bets, with monetary pay-offs, are taken as a *model* of a wider class of phenomena.

Now, Mary did not set out to make a certain loss of £8: we assumed she preferred having more money to having less. Something went wrong. And the only obvious place her rationality can be faulted —the only rationally impermissible part of her reasoning— was that her probability judgements violated the Partition Principle: if she thinks that $A$ is as likely to win as not, and $B$ is as likely to win as not, she must, if rational, think $C$ has no chance of winning at all.

A system of bets which, taken together, yield a certain loss is called (for some reason) a "Dutch Book" and this argument has come to be called the "Dutch Book Argument". My version differs from others in that I do not require the considerations involved in rendering the bets acceptable to be rationally obligatory, only rationally permissible in normal circumstances. Circumstances are, by stipulation, normal. A rationally impermissible result ensues. The only —or at least the most obvious— place to lay the blame is the combination of probability judgements.[11]

There is a second version of the argument. Mary's situation is as before, and she has the same judgements. She is offered a double bet on the first two horses, like the previous ones: she gains £12 if horse $A$ wins, loses £10 if it doesn't; and she gains £12 if horse $B$ wins, loses £10 if it doesn't. She accepts this bet. She is now offered another bet: "You lose £20 if $C$ wins, you gain £2 if $C$ doesn't win."

[11] The locus classicus for this style of justification of the principles of the 'logic of partial belief' is Frank Ramsey's 'Truth and Probability' (1926) in his *The Foundations of Mathematics* (London: Routledge and Kegan Paul, 1931). The argument was discovered independently by Bruno de Finetti. See his 'Foresight: its Logical Laws, its Subjective Sources' (1937) in H.E. Kyburg and H.E. Smokler (eds.), *Studies in Subjective Probability* (New York: Wiley, 1964).

"I certainly do not accept that bet", says Mary. "I think $C$ has a 50–50 chance of winning. I am not prepared to accept a 50% chance of losing £20, together with a 50% chance of gaining only £2. Do you think I'm mad?"

"But you just have, in effect, accepted that bet", we could reply: the double bet she accepted has exactly the same consequences, for each possible outcome, as the single bet she refused: they present, in all relevant respects, the same choice, differently described. In accepting the double bet and rejecting the subsequent offer, Mary is, in effect, saying "Yes please" and "No thank you" to the same thing.

Double bet:

| $A$ | $B$ | $C$ | | $C$ | not $C$ |
|---|---|---|---|---|---|
| +12 | −10 | −10 | equals | −20 | +2 |
| −10 | +12 | −10 | | | |

Again, if we ask what accounts for such blatant irrationality, there is only one plausible candidate: her violation of the Partition Principle. (I leave it to the reader to decide which version of the argument is the more compelling: which is worse, being trapped into saying "Yes" and "No" to the same thing, or being trapped into losing £8 for sure?)

We have considered just one example, but both versions of the argument generalize, provided we can give a general principle for deciding whether odds are acceptable, given one's probability judgements, a principle which is *rationally permissible in normal circumstances*. The principle we need is this. Suppose the probability I assign to a proposition $A$ is $x$. Then a bet "You gain $\alpha$ if $A$ is true, you lose $\beta$ if $A$ is false" is acceptable if and only if $x\alpha > (1-x)\beta$ . [For instance, if I think the probability of $A$ is 1/4, the principle dictates that I accept a bet on the truth of $A$ provided the odds are better than 3:1, for 1/4 times 3 = (1 − 1/4) times

1: I need better than a gain of £3 if $A$ is true, a loss of £1 if $A$ is false, given that I think $A$ has only a probability of 1/4 of being true. If I judge $p(A) = 1/10$, then I need odds better than 9:1 for a bet on $A$; etc.] All we need assume is that this principle is a non-crazy way of deciding whether bets are acceptable in normal circumstances. Then a recipe can be given for constructing systems of bets which entail a certain loss whatever happens when, and only when, one violates the Partition Principle; and for transforming this situation into one where the punter is led to say "Yes" and "No" to the same option differently described.

3.3. Let us call a set of degrees of belief in the members of a partition consistent if and only if they satisfy the Partition Principle. We now have something to say in reply to the sceptic's questions "Why be consistent? What is so good about having consistent sets of beliefs?". By being consistent one avoids the possibility of a certain kind of undesirable consequence of one's actions. But it might be thought that this reply is less than fully adequate, on the grounds that this kind of consequence is rare, and marginal to our practices: the reply, it might be thought, doesn't go to the heart of the matter.

It is instructive to consider the parallel sceptical question about full belief: why should I want my beliefs to be consistent? The obvious reply is that we want our beliefs to be true, and if they are inconsistent, they cannot all be true. The sceptic need not yet give up: she may counter that the elimination of all false belief is an unrealistic aim —we have to live with the fact that some of our beliefs are bound to be false. Why are consistent beliefs, some of which are false, better than inconsistent beliefs, some of which are false? A response might be that the elimination of inconsistencies is part of our methodology for rooting out error and acquiring true beliefs —should we come across an inconsistency, that is a sign that we need to investigate matters further.

Can anything analogous be said about consistency for partial beliefs? It might be thought not: we want our beliefs to be true; the ideal would be to have degree of belief 1 in all truths and 0 in all falsehoods, and no partial beliefs at all. But this is an unattainable ideal. Suppose we had an account of the best attainable degrees of belief to have in given circumstances; a demonstration that such degrees of belief are consistent; and that the elimination of inconsistencies is part of our methodology for improving our degrees of belief —for bringing them closer to the best attainable in the circumstances. This would be a more satisfying defence of consistency —a defence in which consistency is embedded in a richer normative framework. The virtues of consistency alone are not that impressive. The question what makes a set of probability judgements "right", or at least better than another set, in given circumstances, is a very difficult one, which is not addressed here.

### 4. *Consequences of the Partition Principle*

Write "$b(A)$" for a person's degree of belief in $A$ at a time $t$.

(1) $b(\neg A) = 1 - b(A)$

for $\{A, \neg A\}$ is a partition,[12] whose members sum to 1. For example, if $b$(rain tomorrow) = 0.7, $b$(no rain tomorrow) = 0.3.

(2) A contradiction always gets zero degree of belief.[13]

[12] Assuming classical logic applies; see note 9.

[13] As I said above, we are restricting our attention to elementary, recognizable logical facts.

Proof: If $A$ is a contradiction, $\neg A$ is a tautology, so $\neg A$ is certain, so $b(\neg A) = 1$. But $b(\neg A) = 1 - b(A)$. So $b(A) = 0$.

For example, $b$ (It will rain tomorrow and it won't rain tomorrow) = 0.

(3) If $A$ and $B$ are logically equivalent, then $b(A) = b(B)$.

Proof: If $A$ and $B$ are logically equivalent, $\neg A$ and $\neg B$ are logically equivalent, that is, true in exactly the same possible circumstances. As $\{A, \neg A\}$ is a partition, so then is $\{A, \neg B\}$: at least one, and at most one of $\{A, \neg B\}$ can be true. So $b(A) + b(\neg B) = 1$. So $b(A)+1 - b(B) = 1$. So $b(A) - b(B) = 0$. So $b(A) = b(B)$.

For example, $b$ (Either Smith or Jones will lecture tomorrow) = $b$ (It's not the case that neither Smith nor Jones will lecture tomorrow).

(4) If $A$ and $B$ are incompatible, $b(A \vee B) = b(A) + b(B)$.

Proof: Suppose $A$ and $B$ are incompatible. Then, $\{A, B,$ Neither$\}$, i.e., $\{A, B, \neg(A \vee B)\}$, is a partition, whose members sum to 1. But $\{(A \vee B), \neg(A \vee B)\}$ is also a partition, whose members sum to 1. These two partitions have a common member, $\neg(A \vee B)$. So the sum of the remaining members must be the same in each partition, i.e. $b(A \vee B) = b(A) + b(B)$.

For example, {It's square, it's round, it's neither square nor round} is a partition; and {It's either square or round, it's neither square nor round} is another partition. So $b$ (It's either square or round) = $b$ (It's square) + $b$ (It's round).

(5) $b(A) = b(A \& B) + b(A \& \neg B)$.

Proof: $A$ and $((A \& B) \vee (A \& \neg B))$ are logically equivalent. So by (3), $b(A) = b((A \& B) \vee (A \& \neg B)$. But $(A \& B)$

and $(A\&\neg B)$ are incompatible. So, by (4), $b((A\&B) \vee (A\&\neg B)) = b(A\&B) + b(A\&\neg B)$.

For example, $b$ (It will rain tomorrow) = $b$ (It will rain and the temperature will be higher than 10 degrees) + $b$ (It will rain and the temperature will not be higher than 10 degrees).

(6) If $B$ follows logically from $A$ (i.e. $A$ entails $B$), $b(B) \geq b(A)$.

Proof: By (5), $b(A) = b(A\&B) + b(A\&\neg B)$. If $A$ entails $B$, $A\&\neg B$ is impossible, so $b(A\&\neg B) = 0$, so $b(A) = b(A\&B)$.

$$
\begin{aligned}
\text{By (5) again, } b(B) &= b(A\&B) + b(\neg A\&B) \\
&= b(A) + b(\neg A\&B) \\
&\geq b(A).
\end{aligned}
$$

For example, the probability that a thing is square can't be greater than the probability that it has four sides.

(7) $b(A \vee B) = b(A) + b(B) - b(A\&B)$.

This is the general case, when $A$ and $B$ need not be exclusive. For example, the probability of either sun or rain tomorrow equals the probability of sun plus the probability of rain minus the probability of both.

Proof: $\{S \vee R, \neg S\&\neg R\}$ is a partition; so is $\{S\&R, S\&\neg R, \neg S\&R, \neg S\&\neg R\}$ so

$$
\begin{aligned}
b(S \vee R) &= b(S\&R) + b(S\&\neg R) + b(\neg S\&R) \\
&= b(S) + b(\neg S\&R) \\
&= b(S) + b(R) - b(S\&R).
\end{aligned}
$$

From (7) follows the crucial step in our result of section 1: $u(A\&B) \leq u(A) + u(B)$. For $u(A\&B) = b(\neg(A\&B)) =$

$b(\neg A \vee \neg B) = b(\neg A) + b(\neg B) - b(\neg A \& \neg B) = u(A) + u(B) - b(\neg A \& \neg B) \leq u(A) + u(B)$.

These seven results are familiar facts about probability, and follow from the Partition Principle. If it is right to take the Partition Principle as a constraint on a person's degrees of belief at a time, these consequences of it are equally constraints on a person's beliefs at a time.

There are various geometrical, spatial ways of representing probabilities and their interrelations, which many people find easier to grasp than algebraic proofs like the above. There is the Venn Diagram: a rectangle, of area 1, representing the whole range of possibilities, in which we draw circles representing the propositions $A$, $B$, etc., the area of the circle representing the probability of the proposition; the area of the intersection of two circles $A$ and $B$ representing the probability of $A \& B$; etc.

There is the pie diagram: the whole circular pie representing the total range of possibilities; if the size of a wedge-shaped slice represents the probability that $A$, the remainder represents the probability that $\neg A$; the slice representing $A$ can be further subdivided into wedges representing the probabilities of $A \& B$, $A \& \neg B$, etc.

My own preference is for a linear Venn Diagram —or a rectangular pie of length 1, to be sliced in one direction only; for this makes the scale of the probabilities more easily discernible. I hope the role of the following diagrams, as illustrations of the seven results proved above, is clear:

(A comment on the "horizontal spreading" of the later rectangles: all the lines could be superimposed on a single rectangle; they are spread out merely to enable us to see more clearly where the different cuts occur. The abutting rectangles may be thought of as representing different stages in the slicing of the logical cake.)

I once set students the following exercise —as a test of consistency, not of their skills at meteorology. They were to give their opinions about the probabilities of the following occurring in a given city in a particular two-week period in the near future.

(1) It will snow. (2) There will be hail. (3) There will be both snow and hail. (4) There will be either snow or hail (i.e., at least one). (5) Neither snow nor hail. (6) Precipitation but no snow. (7) Snow but no precipitation. (8) Precipitation. (9) Snow or rain or hail or sleet. (10) Snow and the average temperature for the period below normal. (11) Snow and the average temperature for the period at or above normal.

The consequences of the Partition Principle provide a great many constraints on the consistency of answers to these questions. It is not an utterly trivial matter, having consistent degrees of belief (as their answers showed). Still, we have, as yet, a rather small subject —that of consistency of degrees of belief at a time. It becomes much richer once we introduce the notion of a conditional probability; for then we can raise questions about how degrees of belief should change on the acquisition of new information.

It is surprising that these issues are not more discussed in mainstream philosophy of logic. Uncertainty is not a peripheral phenomenon. Yet its proper effect on our reasonings is generally considered to be a subject of optional interest, of which philosophical logicians need not blush to admit ignorance. I have taken just a few beginning steps here. But a lesson of the result of section 1 is that the most obvious generalizations from the special case of certainty are not always the right ones. Uncertainty is a fact of life and we should pay attention to how our reasoning can live with it.

*Recibido: 9 de noviembre de 1995*

RESUMEN

Supongamos que usted tiene un argumento que reconoce como válido y piensa, aunque no con seguridad, que las premisas son verdaderas. ¿De qué manera constriñe esto lo que debe pensar sobre la conclusión del argumento? La primera sección del artículo se aboca a esta cuestión. No se puede sostener que si el argumento es válido y cada premisa es casi segura, la conclusión es casi segura. La Paradoja de la Lotería y la del Prefacio muestran vividamente que esto es erróneo. Muestro que tratando a la falta de seguridad como teniendo la estructura de la probabilidad, la respuesta correcta es ésta: que la falta de seguridad de una proposición sea uno menos su probabilidad. Entonces, si (y sólo si) un argumento es válido, la falta de seguridad de la conclusión no puede exceder la suma de las faltas de seguridad de las premisas. Esto justifica el argüir a partir de premisas casi seguras... en tanto no haya demasiadas. En la Paradoja de la Lotería un número grande de pequeñas faltas de seguridad suman una falta de seguridad máxima.

El resto del artículo es una defensa de la tesis de que la falta de seguridad debe ser tratada en términos de probabilidad (habiendo ya mostrado una atractiva consecuencia de tal tesis). En la sección 2 defiendo que las actitudes epistémicas con respecto a una proposición deben venir en grados para explicar adecuadamente diferencias y similitudes en el comportamiento racional. Y arguyo que la estructura de "grados de creencia" es apta para una idealización mediante una escala numérica. En la sección 3 defiendo el "Principio de la Partición", el cual rige los grados de creencia, y garantiza una estructura probabilística. Llámese "partición" a un conjunto de posibilidades exclusivas y exhaustivas, un conjunto de proposiciones tal que necesariamente una de ellas es verdad pero es imposible que más de una de ellas lo sea. El Principio afirma que los grados de creencia de una persona en los miembros de una partición deben sumar uno (el valor asignado a la seguridad). La sección 4 muestra que el Principio de la Partición, entendido como un requisito de consistencia de grados de creencia, es fértil en consecuencias independientemente plausibles.

Así, este artículo sirve como introducción a una manera fructífera de desarrollar una lógica de la falta de seguridad.

[*Traducción: Raymundo Morado*]