ss

Probabilities and Lewis's Bombshell

[0] Partial Belief.
We invest levels of confidence in propositions. Sometimes we're certain that P, fairly certain that P, uncertain that P, fairly certain that ¬P, certain that ¬P, etc. Our degrees of belief range from full confidence to total mistrust.

Let's measure them with numbers. Let's say full confidence goes with 100, total mistrust goes with 0, and levels in between go with natural numbers in between. Let's say

(*) dob(P) = n $(0 \leq n \leq 100)$.

This is an idealization. But a nice logic for partial belief falls out of it.

[1] The Partition Principle.
Consider a proposition and its negation, say A and ¬A. Think of them as sitting at the ends of a forked path:

A ¬A

@

The actual world, @, is travelling down the path. If it goes left, it's an A-world (a world at which A is true). If it goes right, it's a ¬A-world (a world at which ¬A is true).

Suppose you're trying to reckon whether it'll go left or right. You have evidence pointing to A. So much so, in fact, that you can be 90% sure:

dob(A) = 90.

Picture this with marbles. Think of yourself as having 100 marbles to play with, one for each "degree of certainty". Since you're 90% sure of A, you put 90 marbles by the A-path:

A ¬A

90

@

But where do the other marbles go? What happens to the 10% uncertainty you have about A?

It sticks to ¬A, of course. The other ten marbles go along the ¬A-path:

A	¬A
90	**10**

@

After all: @ will definitely go down one of these paths; and it will definitely go down <u>only</u> one of these paths. So each marble should go along a path. They needn't go along the same path, of course; but they should go along some path or other. This means

$$[dob(A) + dob(\neg A)] = 100.$$

But now consider {A&B,A&¬B, ¬A}. Here too we may picture @ travelling along a forked path:

A&B	A&¬B	¬A

@

And here too logic will guarantee @ goes down exactly one path. Either A is true with B, A is true without B, or A isn't true. Question: how should you distribute your marbles across the paths?

That'll depend on your evidence. But one thing won't: just as before, each marble should be placed on some path or other. For @ will definitely go down one and only one of them. Here too we have

$$[\mathrm{dob}(A\&B) + \mathrm{dob}(A\&\neg B) + \mathrm{dob}(\neg A)] = 100.$$

{A,¬A} and {A&B,A&¬B,¬A} share a special feature. They are both guaranteed by logic to have exactly one true member. Sets like that are called partitions:

(D1) $\{P_1,...,P_n\}$ is a partition $=_{df.}$ logic guarantees exactly one P_i is true.

Think of this just as before. Place the P_is along a forking path:

P_1...........P_2...........P_i...........P_j...........P_n,

@

Since logic guarantees @ goes down exactly one path, each marble should go next to one or another of them. Hence:

$$[dob(P_1) + ... + dob(P_n)] = 100.$$

The principle at work here is the Partition Principle:

(PP) Degrees of belief should sum to 100 across a partition.

[2] Marbles and Venn Diagrams.
There are many ways to represent the distribution of marbles across a partition. A 75/25 split across {A,¬A}, for instance, can be represented by numerals along a forked path (as we've seen); or it can be represented by the area of a "pie" diagram:

or the length of a "chocolate box":

or the area of a Venn Diagram:

But I prefer a mixture of the first and last method. Specifically, I prefer Venn diagrams with numerals:

Figure 1

Here we ignore the area of the A-circle relative to the ¬A containing field. Numerals (rather than area) represent marbles. In other words, numerals represent rational degrees of belief. The "75" inside the A-cirle means your dob(A) equals 75. The "25" outside the A-circle means your dob(¬A) equals 25.

Partitions are easily spotted in these diagrams. Consider, for example, the diagram for A and C:

Figure 2

The area splits into four undivided regions: (I)-(IV). They represent propositions built from A and C:

(I) = A&C
(II) = A&¬C
(III) = ¬A&C
(IV) = ¬A&¬C.

In effect, (I)-(IV) partition the box's area: either A and C are both true, as in region (I); or they're both false, as in region (IV); or one is true while the other false, as in regions (II) and (III). We can think of this in our initial way:

(I) (II) (III) (IV)

@

and we can apply the Partition Principle:

[dob(A&C) + dob(A&¬C) + dob(¬A&C) + dob(¬A&¬C)] = 100.

As we spread marbles across Figure 2, they each fall into some sub-region or other. Any spreading thus amounts to a distribution of confidence across the partition {A&C,A&¬C, ¬A&C,¬A&¬C}.

These diagrams have two nifty features:

--- Propositions are represented by regions of the diagram;

--- Degrees of belief in propositions are represented by numerals within regions.

[3] Marbles, Diagrams and Probabilities.
Suppose A entails C. This means A&¬C is impossible. Let's represent that by shading the A&¬C-region of the Venn diagram:

No marbles go in a shaded region. It represents an impossible state of affairs. Putting a marble there amounts to investing confidence in an impossiblity.

Ok, we're now in a position to see that **degrees of belief, as modelled by these silly little marbles, act like probabilities**. For example, here's a rule of probability theory:

(1) p(¬A) = 1 - p(A).

But the thought here is just this: the number of marbles in the ¬A-region equals 100 minus the number of marbles in the A-region. In other words,

#(¬A) = 100 - #(A).

Since we've got 100 marbles to play with, this is obviously true.

Here's another rule of probability theory:

(2) If A and C are logically equivalent: p(A) = p(C).

But if A and C are logically equivalent, there are no A&¬C-worlds and there are no ¬A&C-worlds. The Venn diagram looks like this:

Since marbles don't go in shaded areas, #(A) obviously equals #(C).

Here's yet another rule of probability theory:

(3) If A and C are incompatible: $p(A \vee C) = [p(A) + p(C)]$.

But if A and C are incompatible, the Venn diagram looks like this:

So #(A v C) equals [#(A) + #(C)].

Here's one final rule:

(4) If A entails C: $p(A) \leq p(C)$.

But when A entails C, the Venn diagram looks like this:

So #(A) is less than or equal to #(C). If there are marbles in the ¬A&C-bit, #(A) is strictly less than #(C). If there are not, they are equal.

[4] Degrees of belief as ratios.
Suppose your spread of marbles across {A,¬A} looks like this:

You've put m marbles in the A-bit, and n marbles in the ¬A-remainder. You're m% sure of A:

dob(A) = m

Notice m is the percentage of marbles in the A-circle relative to the marbles in the entire diagram. Since marbles go either in A or in ¬A, the percentage in A equals the number in A divided by the total number. In other words, dob(A) = m = #(A)/[#(A)+#(¬A)] = m/(m+n).

We can always think of degrees of belief like this. We can always think of them as ratios. They are the number of marbles invested in a proposition divided by the total number we have to invest.

We can also use this ratio idea to understand *suppositional* degrees of belief. We can use it, that is, to understand the degree to which we believe something under the supposition that something else. For consider the Venn diagram for A and C:

Figure 3

The box is split into regions (I)-(IV). Ask this question: given you suppose that A, how confident are you that C?

Well, to suppose that A is to look solely within the A-circle. It's to pretend the A-circle is the entire diagram. In effect, the supposition generates a new Venn diagram:

Figure 4

The circle here corresponds to region (I) of Figure 3; and the rest of the box corresponds to region (II) of that Figure. The entire area of Figure 4 amounts to the A-circle of Figure 3.

This clarifies matters considerably. Your confidence in C under the supposition that A can now be seen as the number of marbles in (I) divided by the total number in Figure 4. It's the ratio of confidence invested in (I) divided by confidence invested across Figure 4:

dob(C/A) = #(III)/[#(II)+#(I)].

In other words,

dob(C/A) = #(A&C)/#(A).

In still other words,

dob(C/A) = dob(A&C)/dob(A).

Putting this in probabilistic terms gives us the fundamental law connecting conditional and unconditional probability:

(CP) p(C/A) = p(A&C)/p(A).

According to this law, conditional probability matches a ratio of unconditional probability.

[5] <u>The Thesis</u>.

A natural thought now presents itself: believing a conditional is believing its consequent while supposing its antecedent. In a phrase: believing a conditional is conditionally believing:

$$dob(A\varnothing C) = dob(C/A) = dob(A\&C)/dob(A).$$

If that's right-- if that's what believing a conditional is-- then the equation will hold no matter how you distribute your marbles. For any distribution of degrees of belief:

$$d(A\varnothing C) = d(C/A) = d(A\&C)/d(A).$$

In probabilisitc terms-- with p ranging over probability distributions-- we have *The Thesis*:

(T) $(\forall p)[p(A\varnothing C) = p(C/A) = p(A\&C)/p(A)].$ [1]

[6] Propositionalism and The Thesis.
Propositionalists claim conditional sentences express propositions. Any two propositions A and C build into a third proposition ø. This is the conditional proposition "running from A to C". In other words:

$$(\forall A,C)(\exists ø)[(A\varnothing C) = ø].$$

Putting this thought together with (T) we get

(!) $(\forall A,C)(\exists ø)(\forall p)[p(A\varnothing C) = p(ø) = p(C/A) = p(A\&C)/p(A)].$

In quasi-English: any two propositions A and C build into a conditional proposition ø so that no matter how probability distributes, the probability of ø will equal the probability of A&C divided by the probability of A. We are led to (!) by two natural thoughts:

(a) Propositionalism about conditionals;

and

(b) The thesis that believing a conditional is conditionally believing.

One of these must go. David Lewis has refuted (!).

[7] The bombshell.

Lewis's proof is hard. It also depends on non-trivial (and occasionally contested) assumptions about the dynamics of rational belief. Fortunately, there's a much better proof to be had.[2]

It goes like this. Take two propositions A and C. Suppose, for reductio, there *is* a proposition ø which makes (!) true. Its truth table will look something like this:

	A	C	ø
1.	T	T	T
2.	T	F	F

3.	F	T	?
4.	F	F	?

The Big Question, of course, centers on 3 and 4: what happens to ø when A is false?

There are three possibilities:

(A) ø is always true when A is false;

(B) ø is never true when A is false;

(C) ø is only sometimes true when A is false;

It's easy to show all three options conflict with (!). Let's take them one at a time.

(A) Suppose ø is true whenever A is false. The truth table is then

	A	C	ø
1.	T	T	T
2.	T	F	F
3.	F	T	T
4.	F	F	T.

Thus:

ø = (A C).

But that conflicts with (!). It's easy to pull apart p(C/A) and p(A C). Just consider the Venn diagram:

As we can see,

p(A C) = [#(I) + #(III) + #(IV)];

while

p(C/A) = #(I)/[#(II) + #(I)].

It's easy to pull these apart. Just distribute marbles this way:

Countless other distributions also do the trick. p(A C) does not equal p(C/A) in all probability distributions. If (A) is the right option, (!) is false.

(B) Suppose ø is never true when A is false. The truth table is then

	A	C	ø
1.	T	T	T
2.	T	F	F
3.	F	T	F
4.	F	F	F.

Thus:

ø = (A&C).

But that conflicts with (!). It's easy to pull apart p(C/A) and p(A&C). Just consider the Venn diagram:

As we can see,

p(A&C) = #(I);

while

$$p(C/A) = \#(I)/[\#(II) + \#(I)].$$

It's easy to pull these apart. Just distribute marbles like this:

Countless other distributions also do the trick. p(A&C) does not equal p(C/A) in all probability distributions. If (B) is the right option, (!) is false.

(C) Suppose ø is only sometimes true when A is false. Its truth table is then not straightforward. Perhaps the truth value of ø in the ¬A-zone varies truth functionally; perhaps it does not. For our purposes it won't matter. No matter how it turns out, the truth table will have this shape

	A	C	ø
1.	T	T	T
2.	T	F	F
3.	F		T

F.

This means the relevant Venn diagram will look like this:

Here the zones are

$$(I) = \neg A \& \neg \emptyset$$
$$(II) = A \& \neg C \& \neg \emptyset$$
$$(III) = A \& C \& \emptyset$$
$$(IV) = \neg A \& \emptyset.$$

(IV) is the ø-bit of the ¬A-region. It corresponds to the upward arrow at line 3 of the previous truth table. Hence

$$p(\emptyset) = [\#(III) + \#(IV)]$$

while

$$p(C/A) = \#(III)/[\#(II) + \#(III)].$$

But it's easy to pull these apart. Just distribute marbles like this:

Countless other distributions also do the trick. p(ø) does not equal p(C/A) in all probability distributions. If (C) is the right option, (!) is false.

Therefore: (!) is false. No proposition's probability matches p(C/A) in all probability distributions. There is just no way a conditional proposition *could* behave in the ¬A-region of space so as to remain aligned with p(C/A). So either

¬(a) Propositionalism about conditionals is false;

or

¬(b) Believing a conditional is not conditionally believing.

---

[1] Here and throughout I assume your confidence in A is not 0.

[2] This proof is a fusion of two others. The first I found in Dorothy's 'On conditionals', MIND 1995. The second I found in Alan Hájek's 'Triviality on the cheap?', Probabilities and Conditionals (CUP:1994), edited by Eells and Skyrms.