Against the Ramsey test

Adam Morton

The Ramsey test is the idea Frank Ramsey first expressed in 1929, that when people accept a conditional 'if A then C' they are 'fixing their degree of belief in C given A'. Later writers have taken the idea to apply only to indicative conditionals, and have interpreted it as saying that, as Allan Gibbard put it, 'the acceptability, assertability, and the like ... of an indicative conditional ... depend upon the corresponding subjective conditional probability'. It is at the heart of Ernest Adams's account of conditionals, and Vann McGee asserts that Adams's account 'describes what English speakers assert and accept with unfailing accuracy'. Take 'depend on' to mean covariance, that the higher the conditional probability the higher the assertability, and call it *Simple Ramsey.*

Jonathan Bennett (2003: 29) develops the idea as follows:

> To evaluate A → C, I should (1) take the set of probabilities that constitutes my present belief system, and add to it a probability = 1 for A; (2) allow this addition to influence the rest of the system in the most natural, conservative, manner; and then (3) see whether what results from this includes a high probability for C.

This is a cautious formulation, since it does not say anything about the degree of assertability of A → C but simply says when you should assert it. We can easily go on to say 'and the strength of the evaluation, the assertability of the conditional, is simply this probability of C'. But we do not have to. We can content ourselves with saying that, for A → C to be assertable, a high probability for C has to come out of the process he describes. Call this *Cautious Ramsey.* (Any Ramsey these days will want to stay on the right side of David Lewis's proofs that conditional probability is not the probability of a conditional – see Milne 2003. The simplest caution here is to restrict the claim to conditionals that do not embed other conditionals.)

I think that both Ramseys are wrong. They embody a misunderstanding of both the function of indicative conditionals and the nature of conditional probability. First I shall give some examples, to get the intuitions out, and develop them before considering objections. Here is a very simple example, to get the general flavour. Suppose I produce a coin you have never seen before and I am proposing to toss it repeatedly. It may come down heads five times in a row. If it does you will suspect it of bias, and

anal_499-513.fm Page 295 Monday, July 19, 2004 4:52 PM

think it very likely that it will come down heads a sixth time. But you do not now think, 'If 5 heads then 6.' You don't even think that if it comes down heads twenty times in a row it will come down a twenty-first, because you know that unbiased coins can come down heads arbitrarily many times. So although Prob(another head|lots of heads) is very high you have no inclination to assert, 'If lots of heads then another head.'

We can get more insight from a pair of more complicated examples, still along classic lines. In the first one we have an array of one hundred and one urns, each of which contains n white balls and $100 - n$ black balls ($0 \leq n \leq 100$). We are particularly interested in urn 1, which contains 99 black balls and one white ball, and urn 0 which has only black balls. A ball is to be chosen from an urn. The conditional probability of the ball's being black given that it was chosen from urn 1 is 0.99. But do we regard the conditional 'If the ball is chosen from urn 1 it will be black' as very assertable? Surely not, because we still believe that the ball could easily be white. And in fact all the conditionals 'If the ball is chosen from urn n it will be black' are intuitively false, with n greater than 0. They all have minimal assertability. It is only when the ball's being chosen from an urn guarantees that it will be black that we regard the conditional as having any credence at all. This is a counter-example to Simple Ramsey. It is also a counter-example to Cautious Ramsey, since we run through the process described and we come up with a high probability for the consequent, but this does not translate into confidence in the conditional.

But the assertability of an indicative conditional does sometimes come in degrees. Consider a second, variant, set-up in which there are 100 urns of which 99 have all black balls and one has all white balls. An urn is chosen at random and then a ball may be taken from the chosen urn. Then though we are not certain of 'If the ball is taken from the urn it will be black' we are pretty confident of it, a confidence that is in rough accord of the conditional probability, again 0.99, of its being black given that it has come from an urn selected from the 100. But of course this probability is also the probability that the urn chosen is an all-black one.

To take these examples in the way I intend, you have to share the intuition that the assertability of, for example, 'if many heads then biased' is less than the corresponding conditional probability. One way of resisting the intuition is to focus on probability rather than assertability. What would we naturally take the probability of 'if 5 heads then 6 heads' to be? It might be suggested that this is just the conditional probability, and that once we grasp this then we will either modify our intuitions about assertability or shift from the confused language of assertion to the clearer one of probability.

It is hard to get unambiguous intuitions about the probabilities of indicative conditionals. I suspect that when you ask people what proba-

anal_499-513.fm Page 296 Monday, July 19, 2004 4:52 PM

bility they think an indicative conditional has they often shift either to the corresponding conditional probability or to the probability of the causal facts underlying the corresponding subjunctive conditional. We can get a rough grasp of the probability we really want by looking for a non-conditional proposition, an 'index proposition', whose probability we are sure is pretty much the same as that of the indicative conditional. For example with the paradigmatic indicative conditional 'If Shakespeare did not write *Hamlet* someone else did', the proposition might be '*Hamlet* is a play that has been around for some time and plays come into existence only when people write them.' For the 'if 5 heads then a 6th' conditional I would suggest as such an index proposition 'the coin is strongly biased (one way or another)', where a strongly biased coin is one that nearly always falls the same way. For given a biased coin we can tell which way it is biased by tossing it a few times. The probability of this proposition will generally be less than the conditional probability of a 6th head given 5 heads. And, more significantly, given a reasonable assignment of prior probabilities as more heads accumulate the proposition's probability given the additional heads will remain lower than the conditional probability.

The reply to this reply is likely to be that it begs the question, by tying the indicative conditional to a non-conditional whose probability transparently does not covary with the conditional probability. My reply can be seen differently, though. It points to a particular sense of the indicative conditional, one that makes tight links with the index propositions suggested. By 'If S didn't write H, someone else did' one can mean something which given one's other beliefs will stand or fall with the facts about H's status as a play and plays in general, and by 'if 5 heads then 6' one can mean something which stands or falls with the coin's being strongly biased. These are permissible senses of the conditionals, which we can distinguish from conditional probability when we see what non-conditional propositions they stand or fall with. (More about distinguishing the relevant sense, below.)

Our intuitions about the assertability of indicative conditionals, understood this way, depend on other factors besides conditional probability. They seem to depend on how the probabilities are obtained, and in particular they do not seem to go along with the averaging-out whereby a random process that operates on the result of another random process can be treated probabilistically in terms of a single probability distribution. Go back to Bennett's formulation. In stage (3) it requires that, after assuming that the probability of A is 1 and letting the effects of this percolate through one's beliefs, one should 'see whether what results from this includes a high probability for C'. Contrast this with an alternative in which instead we see whether what results includes *accepting* C. That is different: the probability of C may be high but below the threshold of acceptance. The lottery paradox has taught us that there is no fixed threshold of acceptance for all propositions. But we could revise the formula to have us see 'whether what results includes a probability for C that is above the threshold for accepting C in the circumstances'.

I think that this revised (3) is on the right track. The break it makes with conditional probability becomes sharper when we consider that the grounds for accepting an indicative conditional often come from considering reasoning that would be set off not by assuming that the antecedent is certain but by assuming that the antecedent is simply accepted. For example a coin has exhibited a run of heads, and we say, 'If it is fair then sooner or later we'll get a tail.' We don't say this because we imagine being certain that it is fair – at least assuming that our evidence is limited to observations of coin tossings this is not something we can be certain about – but because we imagine accepting that it is, and see ourselves forced to expect an eventual break in the run.

We have now revised Bennett's formulation a long way from the initial Ramsey idea, to get: add to our probabilities a probability for A that exceeds the threshold for acceptance, and then see whether this leads to a probability for C that forces us to accept it. Note the 'forces': we don't accept 'if 5 heads then biased' because given the 5 heads we wouldn't be unreasonable to resist concluding that the coin is biased, while we do accept 'if fair then eventually both heads and tails', because it would be unreasonable to expect an infinite unbroken run from a fair coin. The 'conservative' in Bennett's 'natural, conservative' manner of belief-revision is thus highly appropriate, since the indicative conditional has a whiff of epistemic necessity about it, just as the subjunctive conditional has a whiff of causal necessity.

There is another way of putting this. We believe A → C when we think that on learning A we would be in a position to rule out relevant possibilities incompatible with C. (Relevant in the context of utterance.) This fits very naturally on a plausible hypothesis about one function of indicative conditionals. Suppose that we use such conditionals when thinking out in advance how we will expand or repair our belief structures in the eventuality of possible new evidence. We want to lay down a framework that will guide us given a variety of anticipated situations in which we might have to revise our beliefs. But belief revision is a delicate business. On the one hand there is the holistic aspect: how one should change one's beliefs on accepting A depends very delicately on what else one accepts. This leads to the non-monotonic aspect of indicative conditionals, that A → C does not entail (A & B) → C. And this aspect is at least friendly to being interpreted in terms of conditional probability, since in general Prob(C|A) ≠ Prob(C|A & B). On the other hand, we want a certain robustness, and the possibility of piecemeal revision. We do not want our framework to be so sensitive to the overall pattern of belief that in order to know whether, having accepted A → C and then learning that A, one has to search through all of ones beliefs for a pair B, (A & B) → ~C. In response to this second requirement we need a strong connection between antecedent and consequent. In the limiting case we would not commit ourselves in advance to accepting C on learning A unless A logically entails C, but in fact we opt for a weaker strong connection, so that we accept A → C when we can link A to an epistemic context in which there is an ascertainable list of C-defeaters all of which are ruled out by A. On accepting A we have only to know that we are in the right epistemic context to apply A → C, rather than search through all of our beliefs for troublesome Bs. (It seems to me possible that we use the indicative conditional in different ways on different occasions, depending on how much we are pulled by the one or the other of these constraints.)

(You hold a coin that has come down heads 5 times out of 6 and you have to put it in the 'biased' or the 'fair' bin. You think 'probably biased' and then if you have to act then you put it in the biased bin. But if that bin is accumulating too many coins you put it beside the bin because you may revise your judgment later. The indicative conditional I am focusing on is the one that prepares you for future back-tracking, by keeping track of the 'but might still be fair'.)

The difference between the two urn scenarios now makes sense. In the case in which we are simply picking a ball from the urn the fact that the ball has been chosen from an urn that is mostly black does not eliminate the possibility that the lone white ball may have been chosen. In the case in which the ball is picked from an urn which itself has been chosen at random, we might first conclude – reasonably but not inevitably – that the chosen urn was an all-black one. Then in those circumstances, given that conclusion un-inevitable though it was, we are forced to conclude that the ball was black.

The use of indicative conditionals commits one to anticipating belief revision in terms of decisions about what to shift, what to hold on to, and what needs to be ruled out to be sure of a proposition in a context. Conditional probability works in a completely different way. It tries to build lines of future revision into present degrees of (absolute or conditional) belief. It is far from obvious how the two approaches to belief-revision, the traditional one presupposed by our intuitions about indicative conditionals and the probabilistic substitute, relate. Do they produce the same results in the end? Are they deeply incompatible? Do they both fit smoothly into natural human modes of cognition? We don't know the answers to these questions. But I think we do know that they embody very different strategies for anticipating changes in belief.[1]

University of Alberta
Edmonton, AB, T6G 2E5, Canada
adam.morton@ualberta.ca

References

Adams, E. 1975. *The Logic Of Conditionals*. Dordrecht: Reidel.
Bennett, J. 2003. *A Philosophical Guide To Conditionals*. Oxford: Clarendon Press.
Gibbard, A. 1981. Indicative conditionals and conditional probability. In *Ifs*, ed. W. L. Harper, R. Stalnaker, and G. Pearce, 253–56. Dordrecht: Reidel.
McGee, V. 1989. Conditional probability and compounds of conditionals. *Philosophical Review* 98: 485–541.
Milne, P. 2003. The simplest Lewis-style triviality proof yet? *Analysis* 63: 300–303.

[1] This paper was prompted by reading Bennett (2003) and began as a reaction to two pages of that admirable work. Comments by James Hawthorne and Arif Ahmed on a draft forced a lot of re-thinking. Their impact is greatest on the last three paragraphs. The issues touched on in these paragraphs are very hard and I do not pretend to have more than a very tentative grasp of them. If my line is right, though, we can glimpse connections between issues about indicative conditionals, about the concept of knowledge – since the conditional seems to be assertable roughly when if one knows A and then goes on to believe C, which is true, one knows C – and risk-aversion, since the resistance to compounding probabilities embedded in gambles is analogous to Allais' paradox.

Counterfactuals, causal independence and conceptual circularity

JONATHAN SCHAFFER

David Lewis's semantics for counterfactuals remains the standard view. Yet counter-examples have emerged, which suggest a need to invoke causal independence, and thus threaten conceptual circularity. I will review some of these counter-examples (§§1–2), illustrate how causal independence proves useful (§3), and suggest that any resulting circularity is unproblematic (§4).

ANALYSIS 64.4, October 2004, pp. 299–309. © Jonathan Schaffer