

Knowledge, relevant alternatives and missed clues

JONATHAN SCHAFFER

The classic version of the relevant alternatives theory (RAT) identifies knowledge with the elimination of relevant alternatives (Dretske 1981, Stine 1976, Lewis 1996, *inter alia*). I argue that the RAT is trapped by the problem of the missed clue, in which the subject sees but does not appreciate decisive information.

1. Relevant alternatives theory

The RAT identifies knowledge with ‘an evidential state in which *all relevant alternatives* (to what is known) *are eliminated*’ (Dretske 1981: 367). Developing the RAT involves clarifying how to solve for the variables of relevance and elimination in a given case.

Lewis (1996) has developed the most sophisticated version of the RAT to date. According to Lewis, *S* knows that *p* iff *p* is true in every relevant uneliminated possibility. For Lewis, whether or not a possibility is relevant is determined contextually, and he describes seven rules of relevance which clarify how context determines relevance:

Actuality: The possibility that actually obtains is always relevant.

Belief: Any possibility that the subject believes or ought to believe is always relevant.

Resemblance: Any possibility that saliently resembles a relevant possibility is always relevant.

Reliability: Possibilities concerning errors in reliable processes (such as perception, memory, and testimony) are defeasibly irrelevant.

Method: Possibilities concerning errors in sampling and in abduction are defeasibly irrelevant.

Conservatism: Possibilities that are conventionally ignored are defeasibly irrelevant.

Attention: Any possibility explicitly under discussion is always relevant.

These rules are perhaps best regarded as rules of thumb that often enough determine whether a possibility is relevant in a given case.

For Lewis, a possibility w is eliminated for S iff S 's perceptual experience and memory in w would not exactly match his perceptual experience and memory in actuality. For Lewis, perceptual experience and memory are the forms of basic evidence. Lewis's general idea is that the eliminated possibilities are those in which the subject's basic evidence would differ. So, on Lewis's account, S knows that p iff the rules of relevance do not select any possibility in which p is false but S 's basic evidence stays the same.

Alternative versions of the RAT may differ in their views of relevance and/or elimination. I argue that *all* versions of the RAT are trapped by the problem of the missed clue.

2. Missed clues

Missed clues are cases in which the subject sees but does not appreciate decisive information. The subject fails to know despite having conclusive evidence at hand. For instance:

MC: Professor A is testing student S on ornithology. A shows S a goldfinch and asks, 'Goldfinch or canary?' A thought this would be an easy first question: goldfinches have black wings while canaries have yellow wings. S sees that the wings are black (this is the clue) but S does not appreciate that black wings indicate a goldfinch. So S answers, 'I don't know'.

I take it as obvious that S does not know that the bird is a goldfinch in MC. Even S admits it. And Professor A agrees, which is why she fails S. And surely she is right to do so. She would have the support of Barry Stroud: 'I must be able to rule out the possibility that it is a canary if I am to know that it is a goldfinch. Anyone who speaks of knowledge and understands what others say about it will recognize this fact or condition in particular cases.' (1984: 25). Indeed it is hard to imagine any RA theorist thinking otherwise. MC is the RA theorist's *paradigm* of failure of knowledge – a question is explicitly raised and unanswered.

Missed clues are everywhere. Perhaps Sherlock Holmes can follow every clue but we lesser mortals miss clues all the time. The detective may find the fingerprints but fail to match them to the criminal; I may see the landmark but fail to recognize where I am; you may hear the melody but fail to identify the song. The missed clue is like a hieroglyph – one knows its shape but not its meaning.

Missed clue cases are, in a sense, inverted Gettier cases. In a Gettier case, the subject follows evidence that only accidentally points in the right direc-

tion. In a missed clue case, the subject fails to follow evidence that genuinely points in the right direction.

3. *The problem*

The RAT is trapped by missed clues. For Lewis's account to rule that S does not know that the bird is a goldfinch in MC, there must be some relevant and uneliminated possibility in which 'the bird is a goldfinch' is false. This possibility obviously *should* be the canary possibility, since the canary possibility is the only alternative that A raises, and is the basis for the judgement that S lacks knowledge. But Lewis's account rules that the canary possibility is *eliminated*. Since canaries have yellow wings, S's perceptual experience in the canary world would not match his perceptual experience in actuality – in the canary world S would see yellow wings rather than black. The clue would look different.

MC shows that Lewis's account conflates seeing the clue with appreciating what it means. His account treats us all like Sherlock Holmes, able to know where any clue leads.

I anticipate two replies. First, one might look to an account of relevance to generate some further relevant alternative (which of course must turn out uneliminated). Second, one might revise Lewis's definition of elimination to maintain that the canary alternative should count as uneliminated. In responding to these replies I will in effect show that no way of defining relevance and elimination could resolve the problem of the missed clue, which shows that no version of the RAT can escape the missed clue.

4. *Relevance-based responses*

The RA theorist might reply to the problem of the missed clue by looking to the account of relevance to generate some further alternative other than that of a canary. Since the challenge is to capture the judgement that S does not know that the bird is a goldfinch, it is a constraint on any such further alternative that it be uneliminated. One plausible candidate is a canary that looks perfectly goldfinch-like, black wings and all.¹ Perhaps this is a cleverly painted canary, or some sort of genetic mutant (I hereafter call this candidate 'the mutant alternative' and reserve 'the canary alternative' for a normal canary-like canary). In fact I think the mutant alternative is the *only* plausible candidate for a further relevant and uneliminated possibility – in order for the candidate to be relevant it should have to involve some sort of canary, and in order that it be uneliminated it should have to involve something that looks perfectly goldfinch-like.

¹ Lewis (personal communication) suggested this candidate.

The mutant alternative, however, is *irrelevant*. In the context in question, A is testing whether S can discriminate goldfinches from canaries, and so A is *presupposing* that canaries do not look perfectly goldfinch-like (especially with respect to wing coloration). If the mutant alternative were in play, then Professor A would have in effect asked a devious trick-question that not even she herself could answer. The mutant alternative is in effect a *sceptical hypothesis* about bird identification, and easy first questions on ornithology tests do not raise sceptical alternatives.

Here the RA theorist is invited to consider the following pair of cases (variants on a case from Dretske 1970). Case (i): S is at the zoo, watching the striped equines cavort in the cage labelled ‘zebras’. A asks ‘Zebras or mules?’ Here it seems S can easily know that the beasts are zebras, assuming he can decipher the ‘zebras’ sign and/or pick up on the stripes.² Case (ii): as per (i) except that A asks ‘Zebras or cleverly painted mules prankishly placed in the zebra cage?’ In (ii) A has asked a *much* harder question than in (i), and it might seem that here S cannot know that the beasts are zebras. I expect the RA theorist to hold that the difference between (i) and (ii) is that the painted mule question in (ii) raises a sceptical alternative that is not relevant in (i). The RA theorist is then invited to compare the original missed clue case MC where A asks ‘Goldfinch or canary?’ to a variant in which A asks ‘Goldfinch or mutant canary?’ Likewise I expect the RA theorist to hold that the mutant canary question raises a sceptical alternative that is not relevant in the original case.

Lewis’s account of relevance, moreover, correctly rules that the mutant alternative is irrelevant (which is if anything *confirmation* of Lewis’s account of relevance – such a sceptical hypothesis should come out irrelevant in such an innocent context). A brief review of Lewis’s rules of relevance (see §1) shows that only Resemblance and Belief could be thought to render the mutant alternative relevant.³ Now one might think that the mutant alternative is relevant either because it resembles another relevant possibility, actuality, with respect to S’s visual evidence (Resemblance), or because S believes that canaries have black wings (Belief).

Neither Resemblance nor Belief, however, renders the mutant alternative relevant. Starting with Resemblance, while it is true that the mutant alternative resembles actuality with respect to S’s visual evidence, such a respect

² If S cannot decipher the sign or appreciate the significance of the stripes, then one gets a missed clue case.

³ Of Lewis’s other five rules, Actuality just makes the goldfinch possibility relevant, Reliability, Method, and Conservatism all function as limits on relevance, and Attention just makes the canary possibility relevant. So we are left with the goldfinch possibility, the canary possibility, and whatever Resemblance and Belief can add to the mix.

of resemblance cannot count as salient. If such a respect of resemblance were to count as salient, then radical sceptical hypotheses would *always* be relevant (Lewis 1996: 556–7), as would moderately sceptical hypotheses about cleverly painted mules. The salient respects of Resemblance need to be limited in just such a way as to block the mutant from relevance.

Turning to Belief, this rule is really a non-starter, since nothing has yet been said about what S believes.⁴ S might well withhold belief as to whether canaries have black wings (which would be advisable given his ignorance), or S might believe that canaries have red wings while goldfinches have blue wings and so consider either answer equally bad, or S might even believe that canaries have yellow wings but fail to connect his beliefs. In fact, even if S does believe that canaries have black wings (in the dispositional sense of belief), as long as he does not explicitly token this belief at the time, Belief cannot be allowed to apply. For I am disposed to believe that hands can be hallucinated, and many zoo-goers are no doubt disposed to believe that mules can be disguised as zebras. So if Belief were allowed to apply to a subject's dispositional beliefs on a topic, then sceptical hypotheses would *always* be relevant. The applicability of Belief needs to be limited in just such a way as to block the mutant from relevance, in all but the special case where S explicitly tokens the belief that canaries have black wings at the time.

I conclude that the only relevant possibility, or at least the only relevant possibility that stands any hope of being uneliminated, is that of a (normal, yellow-winged) canary. This conclusion should not be surprising, since the canary possibility is just what Professor A was asking about.

5. *Elimination-based responses*

The RA theorist might, alternatively, reply to the problem of the missed clue by looking to the account of elimination to render the (normal, yellow-winged) canary possibility uneliminated. This reply requires some definition of elimination other than Lewis's, since on Lewis's definition the canary possibility is eliminated (§3). Obviously there are many ways of trying to redefine elimination.

Any successful redefinition of elimination, however, will yield the absurd consequence that S could not know that the bird has black wings. In order for the redefinition to succeed in capturing the judgement that S does not know that the bird is a goldfinch, the redefinition must rule the relevant possibility of a normal, yellow-winged canary to be uneliminated. But then,

⁴ Part of why nothing has yet been said about what S believes is that Lewis does not require belief for knowledge. But if one would add a belief requirement, all that would need to be added to MC is that S believes that the bird is a goldfinch, for no good reason.

ipso facto, there would be a relevant and uneliminated possibility in which the bird is yellow-winged. It would then follow from the RAT that S could not know that the bird has black wings.

Now S's belief that the bird has black wings is a true, straightforward perceptual belief, to which (it would seem) no challenges have been raised. S may be ignorant about birds, but he is not colour-blind. In any case, it certainly seems *possible* for S both not to know that the bird is a goldfinch, and to know that the bird has black wings.

Here one might follow-up that S can still know that the bird has black wings, so long as the question of knowledge of wing colour saps the relevance of the canary possibility. On this thought, when A asks 'Goldfinch or canary?' the canary possibility is relevant and uneliminated, but if A were then to ask 'Black wings or yellow?' the canary possibility would no longer be relevant.

But such a follow-up is incompatible with the RA theorist's approach to scepticism, which requires that relevance be not so easily sapped (Lewis 1979: 355). Suppose A raises the brain-in-a-vat scenario, and S asserts that he knows that he has hands. S's assertion rings hollow. But why doesn't S's assertion instead sap the relevance of the vat hypothesis so as to ring true? The standard RA theorist answer is that once a possibility is included as relevant, it will not go away, until perhaps much later when the whole conversation is forgotten. Likewise the canary possibility will not go away just to make S's knowledge of wing colour come out right. In any case, the above follow-up can be blocked, and relevance-fixity achieved by fiat, by asking: 'Is it true in the context generated by A's canary question that S knows that the bird has black wings?' Here the answer is clearly: yes.

I conclude that there is no way of rendering the canary possibility uneliminated without miscounting S as not knowing that the bird has black wings.

6. Conclusion

Missed clue cases show that the RAT cannot account for human knowledge, *however relevance and elimination are analysed*. There are three possibilities worth considering in MC: goldfinch, canary, and mutant. For S not to know that the bird is a goldfinch either the canary alternative must be both relevant and uneliminated, or else the mutant alternative must be both relevant and uneliminated. On Lewis's version of the RAT the canary alternative is eliminated and the mutant alternative is irrelevant. If one revises the RAT to render the mutant possibility relevant and uneliminated (by revising *relevance*), one gets the absurdly sceptical consequence that not even Professor A knows that the bird is a goldfinch, and if one revises

the RAT to render the canary possibility relevant and uneliminated (by revising *elimination*) one gets the absurd consequence that S does not know that the wings are black.⁵

*University of Massachusetts-Amherst
Amherst, MA 01003, USA
schaffer@philos.umass.edu*

References

- Dretske, F. 1970. Epistemic operators. *Journal of Philosophy* 67: 1007–23.
Dretske, F. 1981. The pragmatic dimension of knowledge. *Philosophical Studies* 40: 363–78.
Lewis, D. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8: 339–59.
Lewis, D. 1996. Elusive knowledge. *Australasian Journal of Philosophy* 74: 549–67.
Stine, G. C. 1976. Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies* 29: 249–61.
Stroud, B. 1984. *The Significance of Philosophical Skepticism*. New York: Oxford University Press.

⁵ Thanks to Stewart Cohen, Bredo Johnsen, David Lewis, Jonathan Vogel, an anonymous referee, and especially to Ram Neta.