

EPISTEMIC JUSTIFICATION AND DEDUCTIVE CLOSURE

SAMIR OKASHA

Department of Philosophy, Logic and Scientific Method
London School of Economics

1. My concern here is with the issue of whether knowledge and justified belief are closed under known logical implication. To see what this means, suppose that you know or are justified in believing a proposition p . Suppose that p entails another proposition q , and that you know this entailment to hold. Does it follow that you know or are justified in believing q ? Answer ‘yes’, and you accept the principle that knowledge, or justified belief, is closed under known logical implication. At first glance this principle looks unexceptionable, for valid deduction from known premises seems as good a way as any to extend the scope of one’s knowledge, and likewise for justified belief. But many recent philosophers have rejected closure principles of this sort.¹ One motivation for rejecting closure is that doing so is supposed to help with refuting Cartesian scepticism. The sceptic is construed as arguing from the premises ‘you do not know that you are not a brain-in-a-vat’ and ‘if you are sitting in a chair, you are not a brain-in-a-vat’, to the conclusion ‘you do not know that you are sitting in a chair’, and thus appealing to a closure principle, in contrapositive

¹ Canary and Odergard (1989) provide a good survey of the extensive literature on closure principles. See also Vogel (1990).

form. Those who reject closure concede the sceptic her two premises, but deny that the sceptical conclusion follows. This way with the Cartesian sceptic may look attractive; but I will argue that the reasoning which has led philosophers to reject closure principles is fallacious.

2. Early discussions of closure, inspired by Dretske's and Nozick's writings, tended to focus on knowledge, rather than justified belief. More recent discussions have reversed this trend. I shall follow the example of the recent discussions. Nothing much hangs on this, since closure for knowledge and closure for justified belief amount to the same thing, so long as there are no Gettier-style complications. But justified belief is arguably a more important concept than knowledge for epistemology;² and the examples that have been supposed to refute closure for knowledge would tell equally against closure for justified belief, if they worked at all. So focusing directly on justified belief seems preferable. The principle I shall defend is the following:

(CJ) For all persons and propositions, if a person is justified in believing p , and knows that p entails q , then she is justified in believing q .

3. Two different argumentative strategies have been employed by philosophers who reject closure principles like CJ. The first is simply to describe cases where our intuitions appear to tell us that closure fails. This strategy was employed by Fred Dretske in "Epistemic Operators" (1970), the paper which originated the contemporary closure debate. To take one of Dretske's examples, suppose you are in your local zoo, looking at a group of black-and-

² The view that justified belief is what matters, not knowledge, is especially plausible if scepticism is at issue. Arguably, an interesting form of scepticism must apply not just to knowledge, but to justified belief too. See Williams (1977) for a defence of this view.

white striped animals, enclosed in a cage clearly marked ‘zebras’. Do you know that the animals are zebras? Most people would agree that you certainly do. But do you also know that the animals are not mules, cleverly disguised by the zoo authorities to look like zebras? Dretske’s intuitions tell him that you do not know this. But if the animals are zebras, then they are not cleverly disguised mules, and you know this. So if Dretske’s intuitions are to be trusted, knowledge fails to exhibit closure under known logical implication. You know that the animals are zebras, you know that if the animals are zebras, they are not cleverly disguised mules, but you do not know that they are not cleverly disguised mules. Clearly, Dretske’s example could equally be framed in terms of justified belief. Recently Robert Audi (1988, 1995) has developed a number of similar examples purporting to refute closure principles for both knowledge and justified belief.

The second strategy for denying closure involves developing a general analysis of the concept of knowledge or justification, which has the consequence that closure fails.³ But this second route to denying closure is ultimately parasitic on the first. For the analyses in question are typically defended, or opposed, on the strength of their success, or failure, in accounting for our intuitions in specific cases. Indeed, anyone who does not share the intuitions that Dretske and others appeal to in examples like the zebra case above, will obviously want to reject accounts of knowledge or justification which imply that closure fails. So everything hangs on the plausibility of the alleged counterexamples.

4. In my view there is a simple theoretical reason for *not* rejecting closure principles like CJ. For the concept of epis-

³ Nozick’s (1981) ‘tracking’ account of knowledge falls into this category, as do ‘relevant alternatives’ analyses of knowledge and justification.

temic justification is an essentially normative concept. To call someone's belief justified is to say that the belief is rationally permissible, while to call someone's belief unjustified is to say that it is not. So if the opponent of closure is right, cases exist where it is rationally permissible to believe p but not rationally permissible to believe q , despite knowing that p entails q . Now in practice belief is not an all-or-nothing affair, but comes in degrees. Clearly, believing p must involve assigning p a fairly high degree of belief. So according to the opponent of closure, cases exist where it is rationally permissible to assign p a high degree of belief, but not rationally permissible to assign q a high degree of belief, despite knowing that p entails q . But this conflicts with the Bayesian principle that a rational agent's degrees of belief must conform to the probability calculus. For where p entails q , the Bayesian principle requires the rational agent to assign q at least as high a degree of belief as p ; and if rationality *requires* one to do something, it obviously *permits* one to do it too. Insisting on this requirement in its full generality would no doubt make us all irrational, for some entailments are very complex and hard to spot; but in the case in question, the agent is explicitly stipulated to know that p entails q (for we are dealing with closure under *known* logical implication). So the opponent of the closure principle CJ is committed to a blatant violation of Bayesian rationality constraints.

This argument in favour of closure is by no means conclusive. The posited link between epistemic justification and rationality could be contested, and the Bayesian thesis that probabilistic coherence is necessary for rationality is not entirely uncontroversial.⁴ It is not my aim to delve

⁴ The standard Bayesian defence of the requirement of probabilistic coherence is the 'Dutch-book argument': if your degrees of belief violate the probability calculus, then there exists a set of bets which

further into these matters here. For even if the argument above is accepted, something still needs to be said about the anti-closure intuitions that many philosophers claim to have, before the issue can be considered resolved. A theoretical argument in favour of closure is all very well, but unless bolstered by an explanation of what is going on in the alleged counterexamples to closure, simply generates a paradox. I offer such an explanation below.

5. My hypothesis is this. The alleged counterexamples to CJ are actually counterexamples to a very different, though superficially similar principle, which has been confused with CJ. CJ is true, but this superficially similar principle is false. The intuitions appealed to by Dretske and others certainly exist, but they do not conflict with CJ at all, only with the superficially similar principle.

CJ, to recall, holds that if a person is justified in believing a proposition p , and knows that p entails q , then he is justified in believing q . The principle with which CJ has been confused is the following: if a person *has evidence* for a proposition p , and knows that p entails q , then he has evidence for q . I call this principle ‘CE’. It is easy to see why philosophers may have failed to distinguish CJ from CE, or even deliberately identified the two. For there is surely a close conceptual connection between being justified in believing a proposition, and having evidence for that proposition. Indeed some epistemologists, calling themselves ‘evidentialists’, explicitly hold that these two relations are co-extensional.⁵ I agree that there is a close

you will judge to be fair, but is guaranteed to lose you money. Good expositions and criticisms of the Dutch-book argument can be found in Earman (1992) and Howson and Urbach (1989).

⁵ See Feldman and Conee (1985) for an outline and defence of evidentialism. Williamson (1997) provides a sophisticated defence of the view that the propositions we know are precisely those ones for which we have evidence.

relation between having evidence for a proposition and being justified in believing it. Nonetheless, it is a mistake to identify CE and CJ without qualification. For the concept of evidential support contains a crucial ambiguity, pointed out long ago by Rudolf Carnap, which gives rise to an ambiguity in CE itself. Only on one of its two possible readings, can CE plausibly be identified with CJ. The ambiguity is best brought to light by recalling an old debate between Carnap and Hempel.

6. In “Studies in the Logic of Confirmation” (1945), Hempel tried to provide an explicit theory of the concept of confirmation, or evidential support. The locution on which Hempel focused was ‘evidence e confirms hypothesis h ’. Hempel laid down a number of conditions of adequacy which any acceptable definition of ‘ e confirms h ’ allegedly had to satisfy. One of these was the ‘special consequence condition’ (SCC), which states that if e confirms a hypothesis h , and h entails k , then e also confirms k . In other words, the SCC says that ‘ e confirms h ’ is closed under logical implication. The rationale for the SCC is obvious enough —much of the point of collecting evidence for scientific theories is to raise our confidence in their predictions, a practice which would seem wholly mysterious, if evidential support did not flow down the entailment relation. Hempel’s other conditions of adequacy appeared equally unexceptionable at first sight.

However, in *Logical Foundations of Probability* (1950) Carnap delivered an insightful, and devastating critique of Hempel’s approach (pp. 468–481). According to Carnap, Hempel unintentionally conflated two very different concepts of confirmation. If we say that a hypothesis is confirmed by the evidence, we could mean two different things, Carnap argued. On the one hand, we could mean

that, in the light of the evidence, the hypothesis is better confirmed than it was before. In this sense we might say, for example, that the general theory of relativity is confirmed by gravitational red-shift —the evidence makes the theory more probable than it was beforehand. Alternatively we could mean that, in the light of the evidence, the hypothesis has a high overall degree of confirmation. In this sense, we would probably *not* say that general relativity is confirmed by gravitational red-shift, since most people agree that even after the evidence of red-shift, that theory remains extremely tentative and uncertain.

Following Salmon (1975), I call these two concepts of confirmation the relevance concept and the absolute concept respectively. In probabilistic terms, the difference between the two is clear. In the relevance sense, e confirms h just if $P(h/e) > P(h)$, while in the absolute sense, e confirms h just if $P(h/e) > c$, where c is a constant somewhere between $\frac{1}{2}$ and 1. ('Probability' here means rational degree of belief.) Clearly, e can confirm h in the absolute sense without confirming h in the relevance sense, and vice versa. Carnap then showed that while Hempel usually wrote as if he had the relevance concept in mind, his conditions of adequacy were actually more closely suited to the absolute concept. One of Carnap's results is especially significant: while the SCC does hold for the absolute concept, it does *not* hold for the relevance concept. In other words, if $P(h/e) > c$, and h entails k , then $P(k/e) > c$ too, but if $P(h/e) > P(h)$, and h entails k , it does *not* follow that $P(k/e) > P(k)$.

The ambiguity that Carnap diagnosed in 'e confirms h' equally afflicts the locution 'person X has evidence for h' (and a host of related locutions). In the relevance sense, 'X has evidence for h' means that, in the light of the evidence, X's rational degree of belief in h is higher than it was beforehand. In the absolute sense, 'X has evidence for h' means that, in the light of the evidence, X's ra-

tional degree of belief in h is high. As before, someone can have evidence for a proposition in the relevance sense, without having evidence for it in the absolute sense, and vice versa. Obviously, only in the absolute sense can ‘ X has evidence for p ’ plausibly be equated with ‘ X is justified in believing p ’. In the relevance sense, this equation is totally implausible — evidence that raises the probability of an initially unlikely proposition by a very small amount, does not thereby justify one in believing it. Now consider again the principle CE: if X has evidence for p , and knows that p entails q , then X has evidence for q . Clearly, CE can be read in two different ways, depending on whether ‘has evidence for’ is taken in the absolute or relevance sense. Let us designate these CE_A and CE_R respectively. It follows, therefore, that it is quite plausible to equate CE_A and CJ, but very implausible to equate CE_R and CJ.

7. We are now in a position to explain how anti-closure intuitions arise. The opponent of closure claims to have found an example where we are justified in believing a proposition, but not justified in believing one of its known logical consequences. However, in reality, examples such as these are cases where our *evidence* for a given proposition is not evidence for one of its known logical consequences. *And so long as we are talking about ‘having evidence’ in the relevance sense, this is entirely unmysterious.* But the opponent of closure, like Hempel, fails to notice the crucial ambiguity in ‘ X has evidence for p ’, and thus fails to distinguish CE_A from CE_R . This failure leads him to a radical misappraisal of his own example. In reality, he has provided a counterexample to CE_R . But since he fails to distinguish CE_R from CE_A , and identifies ‘ X is justified in believing p ’ with ‘ X has evidence for p ’, he wrongly thinks he has found a counterexample to CJ, the closure principle for justified belief. But CJ is a very different principle

from CE_R , so producing a counterexample to the latter does nothing to undermine the former. And the fact that counterexamples to CE_R exist is of no particular significance, since CE_R is totally unacceptable anyway.

To support this diagnosis, I shall firstly establish that CE_R really is unacceptable, and secondly, show that the cases presented as counterexamples to CJ really are better understood as cases where CE_R fails.

8. The first task is easy, for CE_R falls to the very same objections that Carnap levelled against Hempel's SCC. Carnap's basic point was simple: if $P(h/e) > P(h)$, and h entails k , then there is nothing in the probability calculus to show that $P(k/e) > P(k)$. Indeed, e can leave unchanged the probability of k , or even lower it. That may sound counterintuitive, but it is simply a fact. Now CE_R holds that, if X has evidence for p , and knows that p entails q , then X has evidence for q , where 'has evidence for' is construed in the relevance sense. Carnap's point applies equally well here: if X assigns $P(p/evidence) > P(p)$, and knows that p entails q , nothing in the probability calculus requires that X set $P(q/evidence) > P(q)$. Indeed, in some cases, X will violate the probability calculus if he sets $P(q/evidence) > P(q)$, and so, arguably, is rationally obliged not to do so. It is simply untrue that evidence for a proposition, in the relevance sense, is automatically evidence for its logical consequences, whether known or not. An example may help to make this clear. Suppose I am drawing cards from a well-shuffled pack. Let p be the proposition that I draw the queen of hearts. Let q be the proposition that I draw a red card. Obviously, p entails q . Now, let e be the proposition that I draw a queen. Clearly, e raises the probability of p from $\frac{1}{52}$ to $\frac{1}{4}$. But equally clearly, e leaves the probability of q unchanged, at $\frac{1}{2}$. Suppose I discover e to be true. Then,

I have evidence (in the relevance sense) for p , but not for its known logical consequence q . Hence CE_R is false.

9. It remains to be shown that the alleged counterexamples to CJ are better understood as cases where CE_R fails. Obviously, examining every alleged counterexample to CJ and similar closure principles is not feasible. Instead, I will focus on two alleged counterexamples from the literature, in the hope that they constitute a representative sample. The first is one alluded to previously: Dretske's case of the zebras in the zoo. To recall, you are in a zoo looking at animals in a cage clearly marked 'zebras'. According to Dretske, you know the animals are zebras, you know that if they are zebras, they are not cleverly disguised mules, but you do not know that they are not cleverly disguised mules. Here is Dretske's reasoning for this opinion:

Do you know that these animals are not mules cleverly disguised by the zoo authorities to look like zebras? If you are tempted to say 'Yes' to this question, think about what reasons you have, what evidence you can produce in favour of this claim. *The evidence you had for thinking them zebras... does not count towards their not being mules cleverly disguised to look like zebras* (1970, p. 1015, my emphasis).

And on this basis, Dretske concludes that you do not know that the animals are not disguised mules, and thus that closure fails. As I noted, the argument works equally for justified belief, if it works at all. But in the light of the foregoing, the flaw in Dretske's argument should be clear. Dretske says that your perceptual evidence for 'the animals are zebras' is not evidence for 'the animals are not cleverly disguised mules'. And in the relevance sense, this claim is certainly plausible. For why should the experience of seeing black-and-white striped animals make you more confident of 'the animals are not cleverly disguised mules' than you

were originally? After all, that experience is just what you would expect if the animals *were* disguised mules. But in the absolute sense, Dretske's claim is not plausible. For as Dretske explicitly admits, your background knowledge gives you very good reason to think that the zoo-keepers have not disguised any mules as zebras. Zoo-keepers do not normally do things like that. So your initial degree of belief in 'the animals are not disguised mules' will be very high, presumably close to one; so your degree of belief in that proposition given the perceptual experience will also be high, unless you take the experience to drastically lower its credibility. But if you do that, you are taking the experience to drastically *increase* the credibility of 'the animals *are* disguised mules';⁶ and it is hard to see why that is a rational thing to do, given your justified belief that you are looking at zebras. It is one thing to say that the perceptual experience does not increase the credibility of 'the animals are not disguised mules'; it is quite another to say that it drastically decreases its credibility. But that is what Dretske must say, if he is to avoid the conclusion that that proposition ends up with a high rational credibility, given the evidence. In short, in the absolute sense the perceptual evidence *is* evidence for 'the animals are not disguised mules', for that proposition has a high probability all along. Only in the relevance sense does the evidential support not penetrate to the known logical consequence.

Therefore, my diagnosis applies neatly to Dretske's example. Dretske moves from 'you do not have evidence for p' to 'you do not know/justifiably believe p', but fails to see the absolute/relevance ambiguity in the former. Only in the absolute sense is the move permissible, but only in the relevance sense is it plausible to say that you do not have

⁶ The point being that $P(\text{not-H}/e) \ll P(\text{not-H})$ is equivalent to $P(H/e) \gg P(H)$.

evidence for p . So correctly described, Dretske's example only shows that having evidence in the relevance sense is not closed under known logical implication, i.e. that CE_R is false. And CE_R certainly is false. But from that it is quite fallacious to conclude that justified belief, or knowledge, is not closed under known logical implication.

A second alleged counterexample to CP comes from Robert Audi's book *Belief, Justification and Knowledge* (1988). On the basis of a number of examples, Audi rejects closure principles for both knowledge and justified belief; the example I will look at happens to be framed in terms of knowledge, but again, would work equally for justified belief, if it worked at all. Audi imagines that he has just heard a vehicle backfiring. This is a sound with which he is well acquainted, well enough to know he has heard a backfire. If the sound he heard was a backfire, then it was not the sound of a firecracker. Audi knows this, and infers that he did not hear a firecracker.⁷ Does he then know that he did not hear a firecracker? Audi thinks it is 'not clear' that he does know this, even though he has validly inferred it from a premise he knows to be true (p. 77).

What leads Audi to this opinion? As with Dretske, Audi's reason appears to be that his perceptual evidence for 'I heard a backfire' is not relevant to 'I did not hear a firecracker'. For Audi admits that from 'general experience' he can conclude that it is improbable that he heard a firecracker (p. 77). But that is insufficient, Audi thinks, to enable him to know or justifiably believe the proposition 'I did not hear a firecracker'. Now this is an odd thing to say —one would have thought that background knowledge, just as much as current perceptual experience, was relevant

⁷ In place of the standard requirement in discussions of closure that the agent knows that p entails q , Audi requires that the agent must have *validly inferred* p from q . This difference does not affect the point under discussion.

to determining the epistemic status of the propositions one believes. The only way of making sense of Audi's remarks, it seems to me, is as follows. What he is really getting at is that his perceptual evidence, in the relevance sense, does not flow down from 'I heard a backfire', to its known consequence 'I did not hear a firecracker'. In other words, the perceptual evidence does not itself increase the rational credibility of 'I did not hear a firecracker', though it does increase the rational credibility of 'I heard a backfire'. That intuition is plausible, and unproblematic —it is another case where CE_R fails. But as with Dretske's example, it tells us nothing about the overall rational credibility of 'I did not hear a firecracker', given the evidence. It does nothing to show that the rational credibility of 'I did not hear a firecracker', given the evidence, is lower than that of 'I heard a backfire', and therefore nothing to show that you are justified in believing the latter proposition but not its known logical consequence. Only a failure to distinguish the absolute and relevance concepts of evidential support could lead one to think that it did.

10. To conclude, nothing I have said demonstrates that justified belief or knowledge is closed under known logical implication. But if the points I have made are sound, they do cast doubt on the reasoning which has led philosophers to think otherwise. They also provide further evidence of the power of probability theory to shed light on issues in general epistemology.

REFERENCES

- Audi, R., 1988, *Belief, Justification, and Knowledge*, Wadsworth, Belmont, CA.
—, 1995, "Deductive Closure, Defeasibility and Scepticism:

- A Reply to Feldman”, *Philosophical Quarterly*, 45, pp. 494–499.
- Canary, C. and D. Odergard, 1989, “Deductive Justification”, *Dialogue*, 18, pp. 305–320.
- Carnap, R., 1950, *Logical Foundations of Probability*, Chicago University Press, Chicago.
- Dretske, F., 1970, “Epistemic Operators”, *Journal of Philosophy*, 67, pp. 1007–1023.
- Earman, J., 1992, *Bayes or Bust: A Critical Introduction to Bayesian Confirmation Theory*, MIT Press, Cambridge, MA.
- Feldman, R. and E. Conee, 1985, “Evidentialism”, *Philosophical Studies*, 48, pp. 15–34.
- Hempel, C., 1945, “Studies in the Logic of Confirmation”, *Mind*, 54, pp. 1–26, 97–121. Reprinted in Hempel, 1965, *Aspects of Scientific Explanation*, Free Press, New York.
- Howson, C. and P. Urbach, 1989, *Scientific Reasoning: The Bayesian Approach*, 2nd edition, Open Court.
- Nozick, R., 1981, *Philosophical Explanations*, Harvard University Press, Cambridge, MA.
- Salmon, W., 1975, “Confirmation and Relevance”, in G. Maxwell and R.M. Anderson (eds.), *Induction, Probability, and Confirmation: Minnesota Studies in the Philosophy of Science*, Vol. 6, University of Minnesota Press, Minneapolis.
- Vogel, J., 1990, “Are There Counterexamples to the Closure Principle?”, in M.D. Roth and G. Ross (eds.), *Doubting: Contemporary Perspectives on Scepticism*, Kluwer, Dordrecht.
- Williams, M., 1977, *Groundless Belief*, Basil Blackwell, Oxford.
- Williamson, T., 1997, “Knowledge as Evidence”, *Mind*, 106, pp. 717–742.

Recibido: 23 de agosto de 1999

RESUMEN

Muchos filósofos sostienen que el conjunto de proposiciones en que creemos justificadamente no se encuentra necesariamente cerrado mediante la implicación lógica conocida. En este artículo propongo que el razonamiento que ha llevado a los filósofos a esta conclusión es falaz. Quienes se oponen a la cerradura confunden el principio de cerradura para las creencias justificadas con un principio de cerradura superficialmente similar, pero muy distinto. Sostengo que el origen de esta confusión se encuentra en el hecho de que los filósofos en cuestión no han advertido una ambigüedad crucial en el concepto de apoyo evidencial señalada por primera vez por Rudolf Carnap.

[Traducción: Héctor Islas A.]