INFERENCE, DEDUCTION, LOGIC[*]

Ian Rumfitt

With characteristic *chutzpah*, Gilbert Ryle tried to carry his 'anti-intellectualist' crusade into the province of logic, long taken to be one of the opposition's strongholds (see especially Ryle 1945, 1946, and 1950). Throughout that crusade, Ryle rightly stressed the importance of our various cognitive skills—skills that cannot be reduced to, or codified as, knowledge of propositions.\[1]/ A good logician will possess many skills of this kind; indeed, what distinguishes a master of the subject is not so much his knowing many logical theorems—*knowledge that*, in Ryle's terms—as his knowing how to deploy them in solving problems. All the same, while Ryle could have contented himself with developing these points so as to amplify and support his claims about the importance and irreducibility of *knowing how to*, he went much further, and offered an account of the nature of logic and of its applicability. That account seems to me to be wrong in almost every particular, but elements of it continue to exert an influence. In this essay, I aim to identify its most basic flaws and to sketch a better treatment of the topic with a view to elucidating the way in which *knowing how to* and *knowing that* interact as we exercise our capacity for deductive argument.

1. Ryle's account of logic

What is logic? According to Ryle, the subject centrally comprises 'formulations of rules of inference or consistency rules' (1946, 236). These 'rules of inference, like the rules of grammar, chess, etiquette and military funerals, are performance-rules' (*op. cit.*, 238). That is to say, they regulate a certain sort of performance: 'references to them are references to criteria according to which performances are characterised as legitimate or illegitimate, correct or incorrect, suitable or unsuitable, etc.' (*ibid.*). Ryle

---

* I am grateful for comments by Jonathan Barnes, Jason Stanley, and David Wiggins.

[1] For a sympathetic elaboration of this aspect of Ryle's thinking, see Wiggins 2009. In his 2011, Wiggins explores the relationship between Ryle's conception of *knowing how to* and Aristotle's conception of the practical in the *Nicomachean Ethics*.

has various terms for the kind of performance which the logician's rules serve to regulate, but the most common (unsurprisingly) is 'inference'. Furthermore, the relevant species of legitimacy or correctness is validity: 'a breach of a rule of logic is a fallacy; an observance of it is a valid inference. To speak of an inference as an observance or as a breach of a rule of logic is only a condensed way of saying that the author of the inference has made his inference in conformity with or in breach of a rule of inference' (*ibid.*).

Ryle distinguishes between two kinds of performance-rule, 'Procrustean rules' and 'canons'. The inference rules of formal logic belong to the first kind, which 'can generally be expressed in brief formulae or terse orders' (*op. cit.*, 240). The 'canonical rules, on the other hand, commonly resist codification' (*ibid.*). The 'principles of induction' (Ryle does not tell us exactly what he takes these to be) are canonical, but there are also canons that bear on purely *a priori* disciplines: 'the formal logician himself in selecting, ordering and proving his Procrustean rules of inference is guided by…non-Procrustean canons'—specifically, by canons similar to those that lead to fertile axiomatizations of mathematical theories (*op. cit.*, 241). Indeed, other canons will have been applied, not in devising the optimal theoretical presentation of the rules of inference, but in discovering those rules in the first place. Aristotle's achievement in this area was to begin to 'crystallize' performance-rules that 'were already being applied' (*op. cit.*, 243). That achievement took a skill of discernment, one which surely cannot be reduced to knowledge of a proposition, but which Ryle compares with the skills displayed by similar 'codifiers', such as Clausewitz (*ibid.*), Izaak Walton and Mrs Beeton (1945, 231). The last comparison inspires a characteristic flourish:

> You couldn't define a good chef as one who cites Mrs Beeton's recipes, for these recipes describe how good chefs cook, and anyhow the excellence of a chef is not in his citing but in his cooking. Similarly skill at arguing is not a readiness to quote Aristotle but the ability to argue validly, and it is just this ability some of the principles applied in which were extracted by Aristotle. Moral imperatives and ought-statements have no place in the lives of saints or complete sinners…Logical rules…are in the same way helpful only to the half-trained (1945, 233).

This somewhat deflating account of the value of logic enables Ryle to dispose briskly of what he regards as the pseudo-problem of explaining how logic applies to the world:

> Some people have worried themselves by speculating how or why the rules of inference apply to the world; they have tried to imagine what an illogical world would be like. But the puzzle is an unreal one. We know already what an illogical man is like; he is the sort of man who commits fallacies, fails to detect the fallacies of others, and so on. The reason why we cannot imagine what an illogical world would be like is that a tendency to flout performance-rules can only be attributed to performers. The world neither observes nor flouts the rules of inference any more than it observes or flouts the rules of bridge, prosody or viticulture. The stars in their courses do not commit or avoid fallacies any more than they revoke or follow suit (1946, 238-9).

'People who construe the logicians' rule-formulae as descriptions of the spine and ribs of the world…assume that a logician's rule-formula "says" something informative. The mistake is not peculiar to them. Other people think that such a rule-formula "says" something uninformative' (*op. cit.*, 243). On a right view, Ryle suggests, these formulae do not say anything at all. When he endorses a rule of inference such as *modus ponens*, the logician is not asserting anything; he is not committing himself to the truth of the claim that the consequent of a conditional invariably follows from that conditional in tandem with its antecedent. Rather, he is formulating a rule with which an instance of inferring can comply or fail to comply, and recommending that our inferences should comply with it. We see here why Ryle's view of logic was grist to his anti-intellectualist mill. If an endorsement of *modus ponens* is not an assertion, then the questions of whether and how one can know the proposition thereby asserted do not arise. So, in this central area of our cognition, we should not be inquiring into the nature of our knowledge of certain propositions. Rather, we should be inquiring into the nature of the rules by conforming to which we can correctly perform certain mental operations.

Ryle's account of logic is elusive. He offers the reader few arguments, but a host of comparisons and similes: all the papers I have cited are littered with occurrences of the expressions 'like' and 'in the same way'. As so often with similes in philosophy, it is hard to know in what respects the things compared are supposed to be alike, or how far to press the putative resemblances.

In what ways are inference-rules like the rules that regulate military funerals? Well, perhaps someone who found himself attending many such ceremonies could discern the operative rules for them, rather as Aristotle discerned some of the rules to which our inferences conform. But the differences between the two cases are more striking than the similarities. The rule that soldiers at a British military funeral carry reversed arms is clearly a local convention; there are countries in which it does not apply. But Ryle, I take it, is not claiming that *modus ponens* is likewise local and conventional. Or, if he is claiming that, he needs to give us an argument. The claim is not to be swallowed on the strength of an unsupported assertion of a likeness.

All the same, those with an interest in the relationship between knowing how to and knowing that will wish to attain an understanding of how different forms of knowledge interact in the case of logical deduction. Ryle mentions Lewis Carroll's famous puzzle in 'What the Tortoise said to Achilles' (Carroll 1895) which he takes to refute the hypothesis that 'knowing how to reason' is 'analysable into the knowledge or supposal of some propositions, namely, (1) the [particular] premisses, (2) the conclusion, plus (3) some extra propositions about the implication of the conclusion by the premisses, etc., etc., *ad infinitum*' (Ryle 1945, 227). That negative conclusion is surely right,\[2]/ and it is also plausible that the codification of logical principles—whether as rules or axioms—must come after we have acquired deductive capacities. But to say this is not to say very much. One wants a detailed account of how those principles emerge from, bear upon, or otherwise relate to our deductive capacities. Ryle does not venture such an account. My task in this essay will be to outline one.

[2] Ryle did not contemplate the reduction of *knowing how to* to *knowing that* proposed by Jason Stanley and Timothy Williamson, according to which knowing how to Φ is a matter of knowing that *this* is a way of Φ-ing, where *this* is presented under a practical mode of presentation (see Stanley and Williamson 2001). Such a view perhaps affords a reply to Ryle's invocation of Carroll: Ryle's slow-witted pupil knows that the conclusion follows from the premisses, but he does not know that *this* (practically presented) way of doing things is a way of deducing the conclusion from the premisses. I am sceptical both about Stanley and Williamson's linguistic arguments for their reductive thesis (see Rumfitt 2003), and about whether the notion of a practical mode of presentation can play the role that they need it to play (see again Wiggins 2009 and 2011). But disputes about the ultimate analysis of *knowing how to*, and about its relationship to *knowing that*, are largely orthogonal to the questions addressed in this paper. Stanley and Williamson will grant that central cases of *knowing how to* differ from central cases of *knowing that* in that the former involve practical modes of presentation of ways of doing things. So an adherent of their position can read the present essay as an attempt to delineate the practical modes of presentation that are involved in making deductions.

2. Inference *versus* deduction

First, though, we need to clear some ground.

As we have seen, Ryle takes the logicians' rules of inference to be performance-rules. Having gone so far, it might seem inevitable that the relevant performances should be inferences. But in fact, if the word 'inference' is taken in its primary sense, the claim that it is inferences that inference-rules regulate is a mistake. The basic problem is that an inference is not a performance: unlike the railway journeys that fascinate Ryle in '"If", "so", and "because"', inferences do not take time, nor are they subject to intentional control. Moreover, *pace* Ryle's position in *The Concept of Mind* (see Ryle 1949, 302-3), an inference is not an achievement of, or an arrival at, a result. An achievement must be something that an agent can try to attain, but it makes no sense to say 'Try to infer "It is either raining or snowing" from "It is raining"'. Alan White got much nearer the mark when he wrote: 'To infer is neither to journey towards, nor to arrive at or be in a certain position; it is to take up, to accept or to change to a position. Inference is not the passage from *A* to *B*, but the taking of *B* as a result of reflection on *A*' (White 1971, 291). At any rate, this captures one focal sense of 'infer', and throughout this essay I shall use the term strictly in this sense.

All the same, there is a species of intellectual activity that the logicians' rules can be thought of as regulating.\[3]/ Sometimes, a thinker engages in the task of tracing out the implications of some premisses. Sometimes, indeed, he does this step by step, taking special care to move only to conclusions that the premisses really imply. Let us call this activity *deduction*. Unlike inferences, deductions do take time, and they are subject to intentional control. They can also be achievements: an examination question might sensibly instruct 'Deduce Gödel's Second Incompleteness Theorem from Löb's Theorem', and a candidate might sensibly report 'I tried to do that but failed'. This sort of intellectual activity is rare in everyday life, but it is central to any discipline—such as mathematics, the sciences, and indeed philosophy—where it is important to draw out the implications of hypotheses in a manner that prevents non-implications from creeping in. Insofar as the term 'inference-rules' suggests that the rules of logic regulate inferences, it is misleading: 'deduction-rules' would have been better.

[3] NB, though: for reasons that will emerge in §§3 and 4, I do not think that the logicians' rules are the *only* rules that regulate deduction.

On this way of understanding the terms, there are many cases where *B* is inferable, but not deducible, from *A*. Indeed, there are cases where *B* is inferable from *A* (but not conversely) while *A* is deducible from *B* (but not conversely). White again: 'We can contrast "From your silence I infer that you have no objections" with "From your lack of objections I deduce that you will remain silent"'.\[4]/ This contrast should occasion no surprise. One often infers *B* from *A* because *B* provides the best explanation of *A*. Thus White's inference is a good one if his colleagues' silence is best explained by the hypothesis that they have no objections to his proposal. But that hypothesis explains their silence in part because one of its implications, in tandem with background facts about White's colleagues, is that they will remain silent.

Can we say anything positive about the relationship between inference and deduction (in the senses specified)? Many philosophers write as though deduction is a species of inference, but on the present understanding of the terms, that must be wrong. Since *dog* is a species of *mammal*, every dog is a mammal, but not every deduction is an inference. Indeed, given that every deduction is a performance while no inference is, *no* deduction is an inference. More interestingly, some deductions do not even issue in an inference. To infer *B* from *A*, we said, is to take up, to accept, *B* as a result of reflecting on *A*. But in drawing out the implications of *A* one may reach *B* without accepting it—and, *a fortiori*, without accepting it as a result of reflecting on *A*. Sometimes a thinker accepts *A* and deduces *B* from it. His acceptance of *B* is then grounded in, or based upon, his acceptance of *A*, and we may describe him as having deductively inferred *B* from *A*. But the deduction of *B* from *A* may not issue in this inference. If *B* is absurd, it may instead issue in the thinker's accepting the negation of *A* on the basis of the negation of *B*. But equally, it may not issue in any inference at all. The thinker's deducing *B* from *A* may make him aware of a relationship between *A* and *B* without leading him to accept, or to reject, either *A* or *B*.

One point these cases bring out is that deduction can play the role we expect it to play in our intellectual economy only if it is applicable in drawing out the implications of false premisses. Ryle, however, entirely overlooks this important point. Possessing a capacity for deduction, he tells us, is 'knowing how to move from acknowledging some *facts* to acknowledging others' (1945, 227; emphasis added). Sometimes he goes further, and writes as though deduction were always a matter of

[4] White *op. cit.*, 292. White holds that ordinary English speakers respect this distinction between 'infer' and 'deduce', a claim which seems to me far-fetched. I claim only that a good philosophy of logic will respect the difference.

drawing out the implications of premisses that we actually know: ‘As a person can have a ticket [for a railway journey from London to Oxford] without actually travelling with it and without ever being in London or getting to Oxford, so a person can have an inference warrant without actually making any inferences and even without ever acquiring the premisses from which to make them’ (Ryle 1950, 250). Ryle does not say what is involved in ‘acquiring’ a premiss, but it appears from the context that knowledge of its truth is required. He would, alas, be far from alone in thinking this. For both Mill and Russell, to deduce is to come to know the conclusion’s truth on the basis of prior knowledge of the premisses. For Frege, a deduction’s premisses must be, if not known, then at least asserted.

Even those philosophers who recognize that we deduce things from false premisses sometimes fail to press the observation as far as it should be pressed. According to Aristotle, whether we are engaged in ‘demonstration’ (i.e., in drawing out the implications of what we know) or in ‘dialectic’ (an enquiry directed towards deciding between two contradictories) ‘makes no difference to the production of a deduction…for both the demonstrator and the dialectician argue deductively after assuming that something does or does not belong to something’—i.e., after assuming that such-and-such is the case (*Prior Analytics* I, 24 a 25-27). That is right, and deduction often assists dialectic by drawing out an absurd or obviously false implication from one of the pair of contradictories between which we are trying to decide, as when the Socrates of the *Theaetetus* draws out absurd implications of the hypothesis that knowledge is perception. But we can also deduce things from premisses that we already know to be false, as when an aged dominie teaching Euclid’s proof that there is no greatest prime number for the fortieth year running begins ‘Suppose there were a greatest prime number, *N*’.

A thinker may, indeed, deduce implications from some premisses whatever his epistemic attitude to them. In order to infer *B* from *A*, one must accept both *A* and *B*. But one may deduce *B* from *A* regardless of whether one knows, believes, wonders about, or disbelieves *A*. This partly explains why the basic criterion of success in deduction is the preservation of truth from premisses to conclusion, rather than the preservation of knowability or assertibility. Consider the argument ‘Suppose Mrs Thatcher was a KGB agent. In that case, she would have taken great care to destroy all the evidence of her treachery. So no one will ever know that she was a Russian agent.’\[5]/ In an appropriate context, that might be a perfectly good deduction, but the conclusion would make no sense if the argument were understood to be elaborating the hypothesis that we know that Mrs Thatcher was a

[5] Frank Jackson and John Skorupski have made cognate points about related conditionals.

KGB agent. In making our deduction, we are drawing out the implications of the truth of the initial supposition, not the implications of our knowing it. To be sure, we sometimes come to know a conclusion by deducing it from premisses that we already know, and in §5 I shall try to explain how we can gain knowledge in this way. But that is not what deduction is. Deduction is a matter of tracing out the implications of premisses, whether we know those premisses or not, and whether they are true or not.\[6]/

3. The varieties of deduction and of implicative relations

The implications of premisses that deduction draws out are not always logical consequences of those premisses. Consider declarer in a game of bridge who reasons in the following way about a line of play, *L*, that he is contemplating using:

> Either East or West has the king of hearts. Suppose that East has it….Then, on that supposition, *L* makes contract. Suppose that West has it.---Then, on that supposition, *L* makes contract. So, either way, *L* makes contract.

The reader is asked to imagine the gaps filled in with detailed deductions about how *L* will play—first under the supposition that East has the king of hearts, then under the supposition that West has it. If the deductions that fill those gaps are sound, then the total deduction will also be sound, and it will owe its soundness partly to its having dilemmatic form. In §4, I shall consider how that formal feature of the total deduction helps to account for its soundness. But the important point for the present is that the deductive capacity being exercised in our argument is specific to bridge. It is the capacity to deduce whether a line of play will make contract under various suppositions about where the unseen cards are. A good bridge player will possess this capacity to a high degree, a poor player to a lesser degree, and a

[6] Some logic text-books—e.g. Lemmon 1965, 8—draw a distinction between premisses and assumptions. Assumptions can be made, or 'introduced', at any stage in the deduction, whilst the premisses are somehow given at the start. But whilst the distinction may help to clarify the way deductions are used in inferences, it is of no relevance to their soundness or validity. The logical rules are applied in just the same way to draw out implications of premisses and assumptions, so we need not dwell on the distinction here.

non-player not at all. It is not a purely logical capacity—although, as our case shows, logical capabilities are involved in it.

Pari passu, the relation of implication, instances of which are traced by exercising this deductive capacity, is not the relation of logical consequence. The implicative relation, too, is specific to bridge. A premiss—such as that East holds the king of hearts—will stand in this relation to a conclusion—such as that L makes contract—if there is no possibility of East's holding the king while L fails to make contract *given that the rules of bridge are adhered to*. The premiss implies the conclusion, one might say, if there is no possibility *within the rules of bridge* that the premiss should be true without the conclusion's being true.

In this simple case, we may define the pertinent implicative relation in terms of logical consequence: some premisses will imply a conclusion if those premisses, together with the laws of bridge, logically entail the conclusion. So it is feasible to employ here what Timothy Smiley has called 'the enthymematic strategy': some premisses X will imply a conclusion B when B follows logically from X together with some further 'tacit' or 'suppressed' premisses Y (see Smiley 1995). Even in this case, though, postulating unexpressed premisses butchers the surface structure of arguments for no good reason, and there will be circumstances where the strategy cannot be applied.\[7]/ Better, then, to think of our intuitive assessments of argumentative soundness as depending on our ability to latch onto the implicative relation that is relevant in the argumentative context. Having done that, we can appraise arguments more or less as they come.

But if we do think about the matter this way, we shall need to say what is characteristic of our deductive capacities, and of the implicative relations that we trace out by exercising them.

Reflection on this problem supports the thesis that implicative relations should have the three 'Tarskian' structural features: they should be reflexive, monotonic, and should manifest the form of transitivity that is captured in the Cut Law.\[8]/ That is to say, where R is any implicative relation, where

---

[7] If a sound deduction from A to B is always to be representable as a logical deduction involving a suppressed premiss, then the logic must license the assertion of a complex premiss $A \rightarrow B$ whenever B is deducible from A. That is to say, the logic must permit the introduction of an operator $\rightarrow$ that validates the deduction theorem. But many logics do not permit this. Only in systems that do will the enthymematic strategy be generally applicable.

[8] Actually, the three features might better be called 'Hertzian' for, as Tarski glancingly acknowledged (Tarski 1930, 62, *n*.1), they had been articulated in earlier publications by Paul Hertz. See Hertz 1922, 1923 and (especially) 1929. For a different—but, I think, complementary—philosophical defence of these three features of implicative relations, see Cartwright 1987.

A and B are individual implicative *relata*,\[9]/ and where X and Y are sets or pluralities of such *relata*, we have:

Reflexivity:	$A\ R\ A$
Montonicity:	If $X\ R\ B$ then $X, A\ R\ B$
Cut:	If $X\ R\ B$ for all B in Y, and $Y\ R\ A$, then $X\ R\ A$.

Reflexivity may be justified as follows. Deduction is a matter of tracing out what our premisses commit us to; in taking A as a premiss we are thereby committed to the truth of A; so where R is any implicative relation—i.e., any relation whose *relata* are traceable by exercising a deductive capacity—we must have $A\ R\ A$. As for the Cut Law, the deductive enterprise of successively elaborating, stage by stage, the commitments of our premisses presupposes that the commitments of those commitments are themselves commitments. At least, this is so if commitments really are *implications* of the initial premisses, as opposed (say) to things that the initial premisses make likely. The Cut Law expresses the relevant form of transitivity, given that many premisses collectively imply a single conclusion. On its face, monotonicity is more doubtful, and in freestyle deductive reasoning, there are many apparent breaches of it. Where the contextually relevant implicative relation is that implicit in Euclidean geometry, the deduction 'This triangle is right-angled; so the square on the largest side is the sum of the squares on the other two sides' is sound. But the deduction 'This triangle is right-angled; the internal angles of triangles sum to more than two right angles; so the square on the largest side is the sum of the squares on the other two sides' is unsound. In this sort of case, though, introducing the additional premiss forces a change in the contextually relevant implicative relation against which the soundness of the deduction is assessed. In a context where it is not assumed that a triangle's internal angles sum to 180 degrees, we cannot apply the Euclidean implicative relation. Rather, we must switch to another implicative relation, one which takes account of further possibilities that are not contemplated in Euclidean geometry. So far from being a counterexample to monotonicity, then, the case illustrates the way in which the operative implicative relation is sensitive to the context of the argument.

[9] Just for the sake of a name, I shall call such implicative *relata* 'statements'. The arguments of this paper do not depend on any theses about the nature of these *relata*. See, though, *n*.12 below.

There is a further, non-structural, property which it is natural to postulate that any implicative relation will possess—that of being *truth-preserving*. By this I mean simply that whenever some premisses imply a conclusion, and those premisses are true, then the conclusion is also true. As with the structural properties, the use we make of deduction presupposes that implicative relations are truth-preserving in this sense. I said earlier that an adequate account of deduction must account for its role in dialectical reasoning. A striking feature of that role is that deducing a conclusion that we know to be untrue from some premisses leads us to infer that at least one of those premisses is untrue. That inference, though, would be unwarranted if the pertinent implicative relation were not truth-preserving. Accordingly, I shall take it to be a further feature of an implicative relation that it is truth-preserving.

In further glossing the notion of an implicative relation, it is indeed natural to do as I did two paragraphs back and invoke a restricted space of possibilities. A stands to B in our bridge-specific implicative relation if there is no possibility within the rules of bridge that A should be true without B's being true; A stands to B in the first of our geometric implicative relations if there is no Euclidean possibility that A should be true without B's being true; and so forth. Quite generally, to each space of possibilities, S, that includes the actual circumstances—i.e., the way things actually are—there corresponds an implicative relation R as follows:

(I) Some premisses $A_1,\ldots, A_n$ R-relate to a conclusion B if and only if, for any possibility x in S, if $A_1,\ldots, A_n$ are all true at x then B is true at x too.

It is easy to verify that a relation defined according to (I) will be reflexive and monotonic, and will obey the Cut Law. Given that the actual circumstances are in S, a relation defined according to (I) will also be truth-preserving.

We often have an antecedent apprehension of a space of possibilities, S, which *via* (I) gives us a grip on the corresponding implicative relation. We may then think of that relation as setting the standard for the exercises of a certain deductive capacity: a deduction in which the capacity is exercised will be unsound if its premisses do not imply its conclusion in the relevant sense. But there is also a converse result which shows how things can work the other way round. That is to say: given a relation, R, that is reflexive and monotonic and obeys the Cut Law, there will exist a space of possibilities, preservation of truth at every member of which is equivalent to R-relatedness. Since, as we saw, exercising a deductive capacity *à outrance* will generate a relation that possesses these three structural features, any deductive capacity is thereby associated with its own characteristic space of possibilities. Furthermore, if R is truth-preserving, then the actual circumstances (i.e., the way things actually are) will be a member of this space.

What guarantees this converse connection is a theorem which Dana Scott attributed to Lindenbaum, but which posterity has insisted upon calling the Lindenbaum-Scott theorem.\[10]/ In one form, the theorem says that whenever we have a non-trivial relation, R, that meets the three Tarskian structural conditions, some premisses will stand in R to a conclusion just in case they stand to that conclusion in every bisective extension of R. Now given such a relation, R, a bisective extension of R is defined by dividing all the statements of the relevant language into two classes, the first of which is closed under R; some premisses then stand in that bisective extension to a conclusion if and only if either the conclusion is in the first class, or one of the premisses is in the second class. We can think of each bisective extension of a relation as describing a possible circumstance as fully as the relevant language permits; a statement will then belong to the first of the two classes defined by the bisection just in case it is true at that possibility. We can also think of the totality of such bisective extensions as corresponding to the totality of possibilities that respect the underlying relation R. On this interpretation, the Lindenbaum-Scott theorem says that some premisses stand in R to a conclusion if and only if the conclusion is true at every one of these possibilities at which all the premisses are true. So the theorem guarantees the existence of a space of possibilities S for which (I) holds. Given also that R is truth-preserving, it follows that the actual circumstances will belong to this space. For this gloss on the Lindenbaum-Scott theorem to be legitimate, the underlying implicative relation must apply between things that are merely supposed to be true, as well between things that really are true. But I argued earlier that we should allow this, a decision that our bridge example confirms. For of the two suppositions made in the course of that argument—that East holds the king of hearts and that West holds it—one must be false.

A judgement that certain premisses imply a conclusion must be distinguished from the deduction of the conclusion from those premisses. And it is the deduction that comes first. By applying one of our deductive capacities, we deduce a conclusion from premisses known or assumed,

---

[10] See proposition 1.3 of Scott 1974. For the version of the theorem employed here, see Koslow 1992, 50-51.

and thereby discover that the premisses stand to the conclusion in the implicative relation that corresponds to that capacity. Deduction is the basic method of discovering instances of implication.

Scott writes of deductions issuing in 'conditional assertions', and some such notion is apposite here. On the strength of a deduction of B from the assumptions $A_1,\ldots, A_n$, we may say: 'Given all of $A_1,\ldots, A_n$, we have B' or 'B, on the assumptions $A_1,\ldots, A_n$'. Some have found the notion of a conditional assertion obscure, but we may understand it by way of a natural generalization of the norms for the speech act of outright assertion. An outright assertion is governed (at least) by the norm of truth: one should not assert B when B is not true. In making an assertion, we present ourselves as conforming to this norm, even if we breach it. A conditional assertion is governed by the corresponding conditional norm: one should not assert B, on the assumptions $A_1,\ldots, A_n$, when all of $A_1,\ldots, A_n$ are true and B is not true. This norm is appropriate for the speech act in which a deduction issues, for the implicative relation that the deduction traces out is assumed to be truth-preserving, so a sound deduction does indeed exclude the case where all the A_i are true and B is not.

Scott took these conditional assertions to be what Gentzen expressed by his sequents or *Sequenzen*. I shall not try to decide how faithfully this reading captures Gentzen's intentions, but let us write 'B, on the assumptions $A_1,\ldots, A_n$' as

$$A_1,\ldots, A_n : B.$$

Pace Scott, though, the colon here is not a sign for a relation. 'East has the king of hearts : L will make contract' means 'L will make contract, on the assumption that East has the king of hearts'; it is a conditional assertion that L will make contract. It does not mean 'The conclusion "L will make contract" is deducible from the premiss "East has the king of hearts"', which is an outright, unconditional assertion that the conclusion stands in a certain logical relationship to the premiss. All the same, when using the colon, some implicative relation is to be taken as understood—from the conversational context or background—as setting the standard for the deduction that issues in the conditional assertion; this relation *will meet the conditions specified earlier.*

Scott's meta-logical reading of the colon permits him to follow Gentzen further, and allow more than one formula to appear after the colon, in the 'succedent' of the sequent. Thus for Scott *the general form of a conditional assertion is* '$A_1,\ldots, A_n : B_1,\ldots, B_m$'. *He understands this to mean:* '*whenever all the statements [$A_1,\ldots, A_n$] are true under a consistent valuation, then at least one [statement in $B_1,\ldots, B_m$] must be true also*' *(Scott 1974, 417; his emphasis). Scott's explanation of* '$A_1,\ldots, A_n : B_1,\ldots, B_m$' *is fine in its own terms. But it endows the explanandum with the sense of a meta-logical statement, one which says that a certain relation obtains between the set (or plurality) of premisses* $A_1,\ldots, A_n$, and the set (or plurality) of conclusions $B_1,\ldots, B_m$. *And to assign such a sense is to change the subject from deductions. In making a deduction, we do not merely identify a finite set of conclusions, one or more of which must be true if all the premisses are true. Rather, we elaborate those premisses—the deduction's initial assumptions—by making specific further assertions within their scope. (While some deductions terminate in a disjunctive conclusion, such a piece of reasoning is adequately represented in the form* $A_1,\ldots, A_n : B_1 \vee\ldots\vee B_m$.) *For this reason, I shall confine the subsequent analysis to conditional assertions with a single statement as succedent.* Although Scott's own proof of the Lindenbaum-Scott theorem was for the multiple-conclusion case, it is straightforwardly adapted to the single-conclusion case (see Koslow, *op. cit.*). So for present purposes we lose no formal power, but maintain the connection with our topic of deduction, by restricting ourselves to succedents with only one member.

4. The role of logic

How does a thinker's specifically *logical* capability relate to these various deductive capacities?

In addressing this question, it helps to begin by comparing the argument given at the start of §3 with the following deduction, which a judge might make in deciding an insurance case:

> Either the deceased committed suicide, or he went skiing alone off-piste. Suppose that he committed suicide….Then, on that supposition, his cover is void. Suppose on the other hand that he went skiing alone off-piste.---Then, on that supposition, he was reckless and his cover is void. So, either way, his cover is void.

Here the gaps are to be filled with legal deductions—that is, with exercises of a deductive capacity which lawyers possess but which non-lawyers lack; the premisses of these subsidiary deductions will include the terms of the relevant insurance policy. This second deductive capacity is quite different from that possessed by a good bridge player (although a single person may possess both), but the two capacities share certain features. In particular, we may think of the legal capacity, too, as answering to a topic-specific implicative relation. Jones's having given Smith £10 in return for Smith's undertaking to deliver certain goods by 1 October, together with Smith's failure to deliver those goods by that date, may be said to imply Smith's liability to compensate Jones for the losses he incurred because of that failure. As before, this is not logical entailment: it is logically possible for Smith to behave in that way without incurring any liability. Rather, it is a relation whose extent is determined by the laws and the precedents of the pertinent jurisdiction.\[11]/

Our two arguments share a dilemmatic form, and owe their soundness in part to their having that form. But how does a thinker's mastery of dilemmatic argument help him to produce sound deductions? The natural—and, I think, correct—answer runs as follows. In each of our cases, the thinker's possession of a certain topic-specific deductive capacity enables him soundly to deduce a conclusion from each of two premisses. His mastery of dilemmatic argument then enables him to splice these two deductions together so as to produce a new sound argument whose premiss is the disjunction of the premisses of its components. The new composite argument is in each case as topic-specific as its parts: in the one case, it is an argument in bridge; in the other, it is a legal deduction. A thinker's logical competence, one might say, consists in an ability to splice together deductions in various fields to produce new, more complex, deductions in those fields. Logical competence, on this view, is a higher-order intellectual capacity: its application yields new deductive capacities from old.

A thinker will possess this higher-order capacity if, in producing new deductive capacities from old, he reliably conforms to certain easily statable rules. And, if the colon of the sequent is understood as in §3, these rules are well formalized as the rules of a sequent calculus—or, more exactly, as the rules of a sequent calculus with single-member succedents. Thus the rule that is applied in both of our dilemmatic arguments may be schematized as follows:

(1) $X, A : C \qquad Y, B : C$

[11] For more about dilemmatic arguments in the law, see Rumfitt 2010b.

$$\overline{\qquad X, Y, A \vee B : C \qquad}$$

Here, A, B and C are arbitrary single formulae, X and Y are arbitrary sets of formulae, and the horizontal line is read as 'so' or 'therefore'. Again, I do not claim that Gentzen had this interpretation in mind when he showed how to formalize classical and intuitionistic logic as sequent calculi (see Gentzen 1935). All the same, a classical logician will accept the classical sequent rules as sound when they are interpreted as general rules for moving from deductions, or conditional assertions, in a given field to other deductions in that field, so long as the implicative relation corresponding to the relevant deductive capacity is held constant throughout the derivation, and so long as it meets our conditions on implicative relations.\[12]/ In particular, nothing in this way of formalizing logic requires that the colon should be taken to signify a notion of specifically logical deduction. On this conception, logical rules are generally applicable rules for forming new deductions from old, not rules that regulate the activity of specifically logical deduction.

This seems to me to be a significant advantage of the account, for it is not obvious in advance of theory what the activity of specifically logical deduction is supposed to be. In particular, it is unclear in advance of theory which implicative relation specifically logical deduction is supposed to be tracing. In the famous passage where he appropriated the word 'entails' from the lawyers to signify the relation of broadly logical consequence, G.E. Moore wrote that we shall 'be able to say truly that "*p* entails *q*" when and only when we are able to say truly that "*q* follows from *p*",…in the sense in which the conclusion of a syllogism in Barbara follows from the two premises, taken as one conjunctive proposition; or in which the proposition "This is coloured" follows from "This is red"' (Moore 1922, 291). Despite the gloss, though, it is hard to be sure what relation Moore had in mind. I know of no logical system that validates the deduction 'This is red; so this is coloured'; and the variety of implicative relations that our ordinary term 'follows' signifies on different occasions of use means that

[12] A classical logician may, though, worry about how completeness is to be secured, given that we are eschewing many-membered succedents. For, in Gentzen's sequent calculus, the operational rules yield intuitionist logic when the system is restricted to single-member succedents; he obtains full classical logic by allowing succedents containing more than one statement. However, we can obtain classical logic with single-member succedents if we take the *relata* of implicative relations to comprise rejections of propositions as false, as well as acceptances of them as true; this is the approach that I recommend to a classical logician in Rumfitt 2000.

a precise apprehension of broadly logical consequence cannot be recovered directly from our understanding of that term.

All the same, even in advance of any delineation of specifically logical deduction, our conception of logical rules makes it easy to see why logic is useful. Being able to deduce conclusions from premisses is clearly useful, if only because it often shows that one or other of those premisses is false. So any thinker will benefit from mastering generally applicable techniques for extending his deductive capacities. On the recommended conception, mastery of the logical rules provides such techniques. In learning to reason about physics, say, a thinker may start with a rather limited deductive capacity. We may pretend, just for simplicity, that his competence in this area is confined to deductions in the form: 'A resultant force is acting on body *a*; so *a* is accelerating'. But if the thinker can reliably contrapose, then his competence will extend to that wider deductive capacity that takes one from the premiss 'Body *a* is not accelerating' to the conclusion 'No resultant force is acting on body *a*'. What is more, mastery of contraposition, and of other logical rules, will also expand his deductive capacities in other fields—indeed, in *all* other fields, given that the logical particles such as 'not' and 'all' are ubiquitous. The theorems of logic may convey no substantive information, but mastery of logical rules expands all a thinker's deductive capacities. Techniques are no less valuable for being applicable only indirectly.

Formula (1) above is a rule, not a statement, so cannot itself be assessed as true or as false. However, its correctness presupposes the truth of a logical law. For rule (1) will be generally applicable in producing correct new deductions from old only if the following law is true:

(2) Whatever implicative relation *R* may be, if *X*, *A* stand in *R* to *C*, and *Y*, *B* stand in *R* to *C*, then *X* and *Y* together with any disjunction of *A* with *B* also stand in *R* to *C*.

Formula (2) expresses the logical law of dilemma, and it illustrates a general thesis: at least in the first instance, logical laws do not characterize some more-or-less elusive relation of specifically logical consequence. Rather, they are general laws governing *all* implicative relations. What is transcendent about the law of dilemma is not that it specially concerns some favoured relation of logical entailment (although, if there is such a relation, the law will apply to it *a fortiori*). Rather, its transcendence lies in its concerning *any* implicative relation, whether it be implication in bridge, in the law, or in anything else. Of course, we are not entitled to assert a general law such as (2) simply on the strength of a couple of favourable cases; apparent counterexamples need to be considered too. In recent discussions, cases involving vagueness and quantum mechanical indeterminacy have been pressed against (2). I cannot discuss these challenges here, but I have tried to show elsewhere how the pressure to restrict law of dilemma when reasoning with vague concepts can be resisted.\[13]/

On the conception of the subject that I am recommending, the basic logical laws will be highly general. Thus, in a sequent calculus in which one and only one formula appears on the right of the colon, the standard rule for introducing the conditional on the left of the colon (i.e., for constructing deductions with a conditional premiss) is:

(3) $$\frac{X, B : C \qquad\qquad Y : A}{X, Y, A \rightarrow B : C}$$

The correctness of rule (3) presupposes the following law:

(4) Whatever implicative relation R may be, if X together with B stands in R to C, and Y stands in R to A, then X and Y together with any conditional whose antecedent is A and whose consequent is B will stand in R to C.

Now in the special case where X is empty, and Y is a singleton whose only member is A, and where C is identical with B, rule (3) reduces to

(5) $$\frac{B : B \qquad\qquad A : A}{A, A \rightarrow B : B}$$

[13] See Rumfitt 2011. See also Rumfitt 2010b, which defends the law of dilemma against a rather different challenge, due to Colin Radford (Radford 1985).

Given that every implicative relation is reflexive, the conditions above the line will be fulfilled no matter what sort of deduction the colon may signify, so the special case reduces further to:

$$(6) \quad \frac{}{A, A \rightarrow B : B}$$

Rule (6) is *modus ponens* and it presupposes the truth of the following law:

> (7) Whatever implication relation R may be, a statement B stands in R to any pair of statements comprising A together with the conditional statement whose antecedent is A and whose consequent is B.

Note that (7)—the traditional logical law of detachment—follows from the more general law (4).

Although I have emphasized the variety of implicative relations which our ordinary deductions trace, the last paragraph points the way to a principled identification of a relation of specifically logical consequence.\[14]/ Law (4) tells us that if certain deductions are sound (by the standards laid down by a given implicative relation) then a related deduction will also be sound (when assessed by the same standards). Some deductions will be sound, though, whatever implicative relation provides the standard for assessing soundness; the conclusion of such a deduction may be said to follow logically from its premisses. From (6), we have that, whatever implicative relation sets the standard for assessing soundness, a deduction by *modus ponens* is sound. So the present account yields the reassuring conclusion that in an instance of *modus ponens* the conclusion follows logically from the premisses. Gentzen's way of formalizing logic has accustomed people to the idea that logical truths are simply the by-products of logical rules—by-products that arise when all the suppositions on which a conclusion rests have been discharged.\[15]/ Our analysis has taken us further in the same direction. On

[14] The relation identified here, though, is 'narrow' logical consequence, not the broader notion invoked by Moore. For an elucidation of that broader notion consonant with the present account of the narrow notion, see Rumfitt 2010a.

[15] Thus Michael Dummett: 'The first to correct this distorted perspective [in which a logic is conceived primarily as a collection of logical truths], and to abandon the analogy between a formalization of logic and an axiomatic theory, was Gentzen…In a sequent calculus or natural deduction formalization of

the conception I am recommending, the classification of deductions as logically valid is itself a by-product of yet more general principles that tell us which deductions stand or fall together when assessed against a given implicative relation. Ascriptions of logical validity are just a limiting case of this wider, relational concern.

5. Knowledge by deduction

Ryle was wrong, I argued earlier, to say that a deduction must start from facts, or known facts. But we sometimes deduce things from premisses that we know, and we value our deductive capacities in part because we can gain knowledge by applying them. So we also need to consider the role that deduction plays in expanding our propositional knowledge.

In some cases, a deductive capacity enables a thinker to gain knowledge that he could not otherwise attain. Suppose I am strapped to the chair in my study. From that chair, I cannot see the street below. I do, however, see that it is raining, and thus know that it is raining. Moreover I know, ultimately on inductive grounds, that if it is raining the street is wet. Accordingly, I reason as follows:

1. It is raining
2. If it is raining, the street is wet

So 3. The street is wet.\[16]/

In this case, exercising my deductive capacity has brought me knowledge that (in my current position) I could not otherwise have attained. In making the deduction, I come to know that the street is wet. *Ex hypothesi*, though, I cannot see the street, so I cannot come to know the conclusion simply by exercising my perceptual capacities, which is how I came to know the first premiss. Similarly, I cannot come to know the conclusion on general inductive grounds, which is how I came to know the second

logic, the recognition of statements as logically true does not occupy a central place…The generation of logical truths is thus reduced to its proper, subsidiary, role, as a by-product, not the core, of logic' (Dummett 1981, 433-4).

[16] We need not worry what the pertinent implicative relation is. For if the rule for introducing → on the left is accepted as regulating the deductive employment of the English conditional, then arguments by *modus ponens* will be sound no matter what the contextually relevant implicative relation may be.

premiss. Even in England, so pessimistic a view of the weather (or of the wastefulness of the water companies) would not yield knowledge. But by exercising my deductive capacity on the knowledge delivered by perception and induction, I can come to know something that I could not know on either of those bases severally.

All the same, cases such as this raise a question. In our example and in others like it, exercising a deductive capacity certainly yields a belief. But under what conditions does belief in a deduction's conclusion qualify as knowledge?

A natural first shot at stating those conditions—a shot, I shall argue, that is rather better than many now suppose—is what we may call the Deduction Principle:

(*DP*) If a thinker knows some premisses, and comes to believe a conclusion by competently deducing it from those premisses, while retaining knowledge of the premisses throughout the deduction, then he knows the conclusion.\[17]/

We clearly need a clause requiring that the thinker should continue to know the premisses: if his knowledge of the premisses were to be destroyed by misleading counter-evidence acquired in the course of making the deduction, then we should not count his belief in the conclusion as knowledge. And we have, I think, enough of a grip on the notion of deductive competence for the Deduction Principle to be more than a tautology. Whatever implicative relation may set the standard for assessing a deduction as sound, some people will be reliable in making deductions only when the premisses really stand to the conclusion in that relation, and others will not. This division gives us our grip on the notion of deductive competence. In fact, in discussing the worries about the Deduction Principle that I wish to address, it will help to focus on the special case of the Principle where the sort of deduction under consideration is specifically logical deduction; and any logic teacher certainly has a grip on the notion of logical deductive competence.

Do we need to add any further conditions to the Deduction Principle to ensure that belief in the conclusion qualifies as knowledge? Perhaps so. Some epistemologists will say that the thinker must not only be deductively competent, but must know, or believe, that he is if his conclusive belief is to qualify as knowledge. Others—more cautiously—will say that he needs not to believe that he is

[17] Compare the formulation of Multi-Premiss Closure' in Hawthorne 2004, 33.

deductively incompetent. Whether one imposes these requirements will depend on one's general epistemological predilections. But two sorts of case have been thought to cast doubt even on the barebones Principle that has been stated. First, there are the so-called 'Dretske cases' of which the following is the most famous (see Dretske 1970). At the zoo one day, you glance into a pen labelled 'zebras' and see a black-and-white horse-like mammal. The animal is, indeed, a zebra, so you know, it seems, that the animal in the pen is a zebra. That premiss entails that the animal in the pen is not a non-zebra carefully disguised to look like a zebra. So, by competently making a (specifically logical) deduction, you come to believe the conclusion—the true conclusion—that the animal is not a non-zebra carefully disguised to look like a zebra. Some philosophers, however, share Dretske's intuition that you do not *know* that conclusion. In order to know it, you would need evidence that excluded the possibility of the animal's being a non-zebra disguised as a zebra, but your inexpert glance into the pen fails to provide such evidence.

A second sort of case involves the accumulation of epistemic risks. A version of the Paradox of the Preface provides a simple example. Suppose you have composed a book comprising only true statements. Suppose too that you know each statement in the book to be true. Now a plausible necessary condition for knowing a statement to be true is that there should be very little risk, given your evidence, that the statement is false. *Ex hypothesi*, then, you meet this condition in respect of each individual statement in his book. Now suppose, however, that you apply the rule of 'and'-introduction to all the statements in the book, thereby reaching a conclusion that is a conjunction of all the individual statements in your book. This seems to be a case of coming to believe a conclusion by competently deducing it (logically) from premisses that you know, so the Deduction Principle tells us that your belief in the conjunction will have the status of knowledge. That claim, though, seems to be inconsistent with the postulated necessary condition for knowledge. For even when the risk of each conjunct's being false is low, the risk of the conjunction's being false will be higher, and if the book contains sufficiently many statements the latter risk can be high enough to disqualify you from knowing the truth of the conjunction, even though the conjunction is true and you believe it.

What lies at the root of this latter objection is a probabilistic conception of epistemic risk. Some philosophers, anxious to ensure that fallible thinkers can acquire knowledge, will wish to say that I can know that it is raining by looking out of the window, even when I am susceptible to occasional hallucinations of rain, so long as the chance of my hallucinating rain is small. Suppose then that I am prone to occasional brainstorms which can do any of three things: they can make me hallucinate rain; they can make me reach inductive conclusions that are not supported by my evidence; and they can make me deduce conclusions from premisses which do not (in the contextually relevant sense) imply them. On the present view, my being susceptible in this way need not preclude me from knowing the premisses (1) and (2) above, or from being deductively competent in the specified sense, so long as the brainstorms have little chance of happening. However, certain events can be individually unlikely without its being unlikely that one of them will happen. On this view, then, an additional condition needs to be met before we can infer that my conclusive belief—my belief that the street is wet—is knowledge. It is not enough that each brainstorm is unlikely to have occurred. There must also be a very low chance that at least one of the brainstorms should have occurred.

How should we react to these cases? Well, we could restrict the original Deduction Principle so as to exclude the apparent counterexamples. (And if the restriction is effected in the manner just suggested, *some* beliefs deduced from known premisses will qualify as knowledge.) On the other hand, we could resist the claim that the cases lately described are counterexamples to the Deduction Principle. They certainly put the Principle under some intuitive pressure, but before accepting them as counterexamples, we should weigh the theoretical costs of restricting the Principle against those of resisting the counterexamples. This in turn suggests is that we should consider what the ground of the Deduction Principle might be. Once identified, that ground ought to show what restrictions, if any, the Principle needs.

A first shot at grounding the Deduction Principle might run like this. If a thinker qualifies as deductively competent (in a given argumentative context), then he will be disposed to deduce a conclusion from some premisses only when the conclusion really does stand in R to them, where R is the implicative relation that sets the standard for assessing deductions in that context. Now suppose that a deductively competent person knows the premisses of an argument, and deduces its conclusion from those premisses. Because he knows the premisses, those premisses are true. And because the premisses are true, and the conclusion is R-related to them, the conclusion is also true. (Since R is a implicative relation, it will preserve truth from premisses to conclusion.) *Ex hypothesi*, our thinker is deductively competent, so he will deduce a conclusion from some premisses only if they really imply the conclusion (in the contextually relevant sense of 'imply'). Accordingly, when a deductively competent thinker deduces a conclusion from premisses that he knows, the belief thereby formed will be true, and it will have been produced in a way that reliably yields true beliefs. Suppose finally that we accept a reliabilist conception of knowledge. On that conception, what endows a true belief with the status of knowledge is precisely that it has been produced in a way that reliably yields true beliefs. So, assuming a reliabilist conception of knowledge, belief in a conclusion that has been competently deduced from known premisses will have the status of knowledge, so that the Deduction Principle is unrestrictedly true.

As we shall see, the proponent of this argument puts his finger on something important when he focuses on the connection between deductive competence and implication. But the argument as it stands is vulnerable to an objection, even if we accept the reliabilism needed at the last step. The objection is that the putative explanation of the truth of the Deduction Principle cannot be right, for if correct it would explain too much. In the explanation, the only use that is made of the hypothesis that the thinker knows the premisses of his deduction is to derive the claim that those premisses are true. So the same account would appear to explain the truth of the Pseudo-Deduction Principle:

(*PDP*) If a thinker comes to believe a conclusion by competently deducing it from true premisses that he believes, then he knows the conclusion.

But the Pseudo-Deduction Principle is patently absurd. Coming to believe a conclusion by competently deducing it from some true beliefs reliably—indeed, infallibly—yields a true belief. But it clearly does not constitute a reliable way of producing true beliefs in the sense needed for a belief that is thereby produced to qualify as knowledge. For (*PDP*) does not require that the thinker should know the premisses of his deduction.

I think we can find an improved ground for the Deduction Principle that escapes the objection if we reflect on what Mark Sainsbury has called the 'Reliability Conditional':

(*RC*) If a thinker knows that *P*, then his belief that *P* could not easily have been wrong.

The modality here relates to the knowing, rather than to what is known. To put the point in terms of possible worlds, if a thinker knows that *P*, there is no nearby world (no world which could easily have been actual) in which he falsely believes that *P*—or, better, in which he falsely believes that *P* on the same basis as he actually believes that *P*. Although Sainsbury himself advances (*RC*) only as a necessary condition for knowledge, those with reliabilist sympathies in epistemology will be tempted to take it to be sufficient as well: on this view, a true belief that is formed in such a way that it could not easily be wrong will have the status of knowledge.

How does the Reliability Conditional bear on knowledge acquired by deduction? To see how it does, let us return to our original deduction and consider how I could have believed falsely that the street is wet on the same basis as my actual belief (which is true). Now the actual basis of that belief is a deduction in which the bases for my premisses are spliced together to form a basis for my conclusion. And we know that if true premisses imply a conclusion, then that conclusion will be true. So in any nearby world in which my belief that the street was wet has its actual basis but is false, at least one of the following three conditions must be met:

(1) my belief that it is raining has its actual basis, but is false;

or

(2) my belief that the street is wet if it is raining has its actual basis, but is false;

or

(3) I deduce the conclusion from my premisses, but in fact my premisses do not imply my conclusion.

Now we are supposing that I know the first premiss of my argument. By the reliability conditional, then, there is no nearby world in which possibility (1) obtains. Similarly, given that I know the second premiss, there is no nearby world in which possibility (2) obtains. Finally, it is a mark of deductive competence that when a thinker deduces a conclusion from some premisses, it could not easily have been the case that his premisses fail to imply his conclusion. Then, given that I am deductively competent, and that *all nearby worlds belong to the space of possibilities associated with the relevant implication relation*, there is no nearby world in which possibility (3) obtains either. There will be cases where the italicised condition does not obtain. However, when we are dealing with logical deduction, as in the special case of (*DP*) that we are considering, the condition will obtain: for any way in which things could easily have been is logically possible. But we said that, in any nearby world in which my belief that the street was wet has its actual basis but is false, at least one of our three conditions must be met. And we have just argued that, when the conclusion is competently deduced from known premisses, none of these conditions is met in any nearby world. Accordingly, there is no nearby world in which the conclusive belief has its actual basis but is false. So, if we accept the converse of (*RC*), that conclusive belief will qualify as knowledge. Thus (*RC*) in tandem with its converse vindicate the Deduction Principle.

This analysis brings out the special role of deduction in a way that explains why this vindication of the Deduction Principle does not extend to vindicate the Pseudo-Deduction Principle. The 'method' of forming beliefs that consists in only believing what follows from true beliefs is reliable, indeed infallible, for a false belief cannot follow from true beliefs. But that 'method' does not deserve the title, for applying it presumes some antecedent criterion for judging whether the premisses are true. What a deductive capacity provides, then, is not a new method for forming beliefs *per se*, but rather a means of combining reliable methods of belief-formation that one already possesses so as to yield a new *method* which has a wider range of application than its components. A deductive capacity, one might say, yields a second-order method of belief formation. This method may itself be applied to establish the truth of certain statements—namely, those implied by all premisses (or none) under the relevant implicative relation. But that is not the present case, nor the central one. Instead, the value of such capacities lies in their power to splice together reliable methods of belief-formation so as to yield further reliable methods which have a wider range of application than their components. Once this is clear, it will be clear why our ground for the Deduction Principle does not extend to justify the Pseudo-Deduction Principle.

What, though, of the apparent counterexamples to the Deduction Principle? Our analysis gives us the resources, I think, to resist both of them. In the Dretske example, there are two cases to consider. Either there is a joker at the zoo who is disposed to disguise non-zebras as zebras—in which case the subject's belief that the animal in the pen is a zebra could easily have been wrong, so that he does not know the deduction's premiss; or there is no such joker—in which case the subject's belief that the animal in the pen is not a non-zebra disguised to look like a zebra could not easily have been wrong, so we may allow that he knows the conclusion. Either way, the case poses no threat to the Deduction Principle. Matters are similar with the Paradox of the Preface. At any possible world where the long conjunction is false, at least one conjunct is false. Now if the long conjunction could easily have been false, there is a nearby possible world at which it is false. At that nearby world, though, at least one of the conjuncts will be false, showing that one of the conjuncts could easily have been false. That is to say, if the long conjunction could easily have been false, then the author does not, after all, know every statement in his book. So either he does know every statement in the book, in which case the long conjunction could not easily have been false, so that we may credit him with knowledge of the conjunction; or his book contains a statement that he does not know, in which case the Deduction Principle is inapplicable. Properly analysed, then, the Paradox of the Preface also provides no counterexample to the Principle.

My endorsement of the Deduction Principle is tentative. Perhaps there are other examples which expose a flaw in the justification advanced for it, and show how the Principle needs to be restricted. Further exploration of the issue must be left to the epistemologists, but it is interesting that we have a justification of the Principle that depends only on what I take to be an attractive general theory of deduction and the plausible epistemological thesis (*RC*). Surprisingly many contemporary epistemologists are willing to reject the Principle. Our analysis at least reveals the high cost of doing so.

6. Conclusion

How does the proffered account of deduction bear on Ryle's general picture of the mind?

At the heart of that picture is the attack on 'intellectualism', the view that what marks out intelligent behaviour is its being 'piloted by the intellectual grasp of true propositions' (Ryle 1949, 26). And, I think, our analysis of deduction can contribute to the Rylean enterprise of subverting that picture. One might put the point this way. On an intellectualist view, the paradigm of intelligent behaviour is theory construction, the goal of which 'is the knowledge of true propositions or facts' (*ibid.*). As we have seen, deduction plays an important role in theory construction—both in refuting false hypotheses and in further elaborating what we have come to know. However, deduction itself is not piloted by the knowledge of true propositions. A deductive capacity cannot consist in knowledge of true propositions—whether those propositions are logical truths or propositions to the effect that this statement follows from these others—for a thinker could know the propositions while lacking any capacity to make deductions. That is the moral of Lewis Carroll's fable. Rather, a deductive capacity is an intellectual ability, exercises of which can (among other things) extend our knowledge of premisses to yield further knowledge of what those premisses imply. As we have seen, such an ability can be turned on itself to yield knowledge of logical truths and knowledge of instances of the relevant implicative relation. But that knowledge is a by-product of a deductive capacity: it is not what pilots the capacity. A thinker could have a deductive capacity even if he had never turned it on itself to attain knowledge of truths implied (in the relevant sense) by all premisses or none. And he could have such a capacity even if he lacked the notion of consequence or implication. So, at the very basis of the intellectualist theory of the mind is a form of intelligent behaviour—namely, the deducing of the implications of premisses—which the theory cannot accommodate.

Is a deductive capacity a form of Rylean *knowing how to*? The answer seems unclear. In some respects, it resembles Ryle's paradigms of that kind of knowledge: it is an ability which one acquires—indeed learns—by practice, and which one can improve by drill (or so those of us who are called upon to teach elementary logic must hope). On other hand, the very fact that it is possible to state deduction-rules puts it at the other end of the spectrum from such uncodifiable instances of Rylean savoir faire as the abilities to prosecute a war, to fish, or even to cook. But if the answer is unclear, so is the significance of the question. It matters greatly to Ryle's anti-intellectualist crusade that some intelligent behaviour should not be piloted by knowledge that. But, so far as I can see, it matters not a jot that all the intelligent behaviour that is not so piloted should be a manifestation of the agent's knowing how to do something.

Without compromising his wider philosophical project, then, Ryle could have said what I have said about deduction and logical rules. I can only hope what I have said is closer to the truth than what he actually said.

REFERENCES

Carroll, L. 1895. ‘What the Tortoise said to Achilles’. *Mind* **4**: 278-80.

Cartwright, R.L. 1987. ‘Implications and entailments. In his *Philosophical Essays* (Cambridge, Mass.: MIT Press), pp.237-56.

Dretske, F.I. 1970. ‘Epistemic operators’. *The Journal of Philosophy* **67**: 1007-23.

Dummett, M.A.E. 1981. *Frege: Philosophy of Language*, 2nd edition. London: Duckworth.

Gentzen, G. 1935. ‘Untersuchungen über das logische Schliessen’. *Mathematische Zeitschrift* **39**: 176-210, 405-431. Translated as ‘Investigations into logical deduction’ in *The Collected Papers of Gerhard Gentzen*, ed. and trans. M.E. Szabo (Amsterdam: North Holland, 1969), pp.68-131.

Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Clarendon Press.

Hertz, P. 1922. ‘Über Axiomensysteme für beliebige Satzsysteme I’. *Mathematische Annalen* **87**: 246-69.

——. 1923. ‘Über Axiomensysteme für beliebige Satzsysteme II’. *Mathematische Annalen* **89**: 76-102.

——. 1929. ‘Über Axiomensysteme für beliebige Satzsysteme’. *Mathematische Annalen* **101**: 457-514.

Koslow, A. 1992. *A Structuralist Theory of Logic*. Cambridge: Cambridge University Press.

Moore, G.E. 1922. *Philosophical Studies*. London: Routledge and Kegan Paul.

Radford, C. 1985. ‘The umpire’s dilemma’. *Analysis* **45**: 109-111.

Rumfitt, I. 2000. '"Yes" and "No"'. *Mind* **109**: 781-823.

——. 2003. 'Savoir faire'. *The Journal of Philosophy* **100**: 158-66.

——. 2010a. 'Logical necessity'. In R.L.V. Hale and A. Hoffmann, eds., *Modality: Metaphysics, Logic, and Epistemology* (Oxford: Clarendon Press), pp.35-64.

——. 2010b. 'Ricky Ponting and the judges'. *Analysis* **70**: 205-10.

——. 2011. 'The logic of boundaryless concepts'. Forthcoming in Alexander Miller, ed., *Logic, Language, and Mathematics: a Festschrift for Crispin Wright* (Oxford: Clarendon Press, 2011).

Ryle, G. 1945. 'Knowing how and knowing that'. *Proceedings of the Aristotelian Society* **46**: 1-16. References to the reprint in Ryle 2009, pp.222-35.

——. 1946. 'Why are the calculuses of logic and arithmetic applicable to reality?' *Proceedings of the Aristotelian Society, Supplementary Volumes* **20**: 20-29. References to the reprint in Ryle 2009, pp.236-43.

——. 1949. *The Concept of Mind*. London: Hutchinson.

——. 1950. '"If", "so", and "because"'. In Max Black, ed., *Philosophical Analysis: A Collection of Essays* (Englewood Cliffs, NJ: Prentice Hall) pp.302-318. References to the reprint in Ryle 2009, pp.244-60.

——. 2009. *Collected Papers, volume 2*. London: Routledge.

Sainsbury, R.M. 1997. 'Easy possibilities'. *Philosophy and Phenomenonological Research* **57**: *907-19.*

Scott, D.S. 1974. 'Completeness and axiomatizability in many-valued logic'. In L. Henkin et al, eds., *Proceedings of the Tarski Symposium* (Providence, RI: American Mathematical Society), pp.411-35.

Smiley, T.J. 1995. 'A tale of two tortoises'. *Mind* **104**: 725-36.

Stanley, J.C. and T. Williamson. 2001. 'Knowing how'. *The Journal of Philosophy* **98**: 411-44.

Tarski, A. 1930. 'Fundamentale Begriffe der Methodologie der deduktiven Wissenschaften I'. *Monatshefte für Mathematik und Physik* **37**: 361-404. Translated by J.H. Woodger as 'Fundamental concepts of the methodology of the deductive sciences', in Tarski 1983, pp.60-109.

——. 1983. *Logic, Semantics, Metamathematics*, eds. J.H.Woodger and J. Corcoran. Indianapolis: Hackett.

White, A.R. 1971. 'Inference'. *The Philosophical Quarterly* **21**: 289-302.

Wiggins, D.R.P. 2009. 'Knowing how to and knowing that'. In Hans-Johann Glock and John Hyman, eds., *Wittgenstein and Analytic Philosophy: Essays for P.M.S. Hacker* (Oxford: Clarendon Press), pp.263-77.

——. 2011. 'Practical knowing'. Forthcoming in *Mind*.