STEPHEN READ

FORMAL AND MATERIAL CONSEQUENCE

1. INVALIDITY

How do we show that an argument is invalid? Consider this example:

(1) All cats are animals
Some animals have tails
So some cats have tails.

The premises are true and so is the conclusion. Yet there is an obvious sense in which the truth of the premises does not guarantee that of the conclusion. The argument is invalid. But how can we show that invalidity?

One thought is, that arguments of the same sort, or form, actually lead from truth to falsity. Although their premises are true, their conclusions are false. The same could have been true of (1), though in fact it isn't. What we might try, therefore, is to formalise the argument, and show that the form is invalid. Using the above example, we obtain

(2) $(\forall x)(Fx \rightarrow Gx), (\exists x)(Gx \,\&\, Hx) \vdash (\exists x)(Fx \,\&\, Hx)$,

where Fx reads 'x is a cat', Gx is 'x is an animal' and Hx is 'x has a tail'. To show this form is invalid, we find another instance of it, with a different key, but in which, though the premises are still true, the conclusion is false. For example, we might let Fx and Gx read as before, but let Hx read 'x is a dog':

(3) All cats are animals
Some animals are dogs
So some cats are dogs.

The premises are true and the conclusion false. So this argument really is invalid. Since every instance of a valid form is valid, (2) is an invalid form.

Journal of Philosophical Logic **23**: 247–265, 1994.

Does this show that (1) is invalid? Not immediately; for every valid argument is in instance of some invalid form. For example, every two premise argument is an instance of the form

(4) $P, Q, \vdash R,$

which is patently invalid. But that does not show that every two-premise argument is invalid.

The trouble with (4), of course, is that it does not reveal sufficient structure. What we try to do when we formalise an argument like (1) is to articulate its structure so that if there is a dependency of the conclusion on the premises it will be revealed. Such a technique is ideal when we find a form which is valid and of which the argument is an instance. But what if we cannot – as with (1)?

What we may be tempted to say is that (2) reveals as much structure in (1) as can be revealed. Since (2) is invalid, this means that (1) has failed its best possible chance to be shown valid, and so must be invalid.

It does show that (1) is not valid in virtue of its form. But does that show it is not valid? How else might it be valid? In a valid argument, the truth of the premises must somehow rule out the falsity of the conclusion. So it must be impossible for the premises to be true and the conclusion false. Could the premises of (1) be true and the conclusion false?

Suppose the world were much as it is now, but cats evolved to become tailless – Manx cats take over, say. In such a world, all cats are animals, some animals have tails (cats no longer do, but dogs are unchanged), but now no cats have tails. We have represented to ourselves a situation in which the premises are true and the conclusion false. So the truth of the premises does not guarantee that of the conclusion. (1) is invalid.

John Etchemendy (1990) contrasts "interpretational" with "representational semantics". In representational semantics we describe a situation, perhaps different from how things actually are, in which the propositions take various values. In interpretational semantics, we interpret certain expressions differently from their actual interpretation to much the same effect. When we formalised (1) as (2) and then interpreted the predicate letters in (2) to obtain (3), we varied the interpretation – we effectively interpreted 'have tails' to mean 'are dogs'. Under this interpretation (retaining the standard interpretation of 'cat' and 'animal' but varying that of 'have tails') the premises come out true and the conclusion false. Similarly, representing a situation in which no cat has a tail, we describe one in which the premises of (1) are true and the conclusion false. Either way, whether interpretationally or representationally, we seek to show that (1) is invalid.

The terminology of 'interpretation' versus 'representation' is not altogether a felicitous one. Talk of representations suggests unwanted mental models, representing things to ourselves; and in both cases, interpretations are what is at stake, on the one hand, keeping interpretations fixed but considering varying situations, on the other, keeping the situation fixed by allowing the interpretation to vary. A more perspicuous terminology might be to speak of substitutional semantics on the one hand, where we substitute different expressions within a substitution-class for certain expressions, to see if truth results;[1] and of modal semantics on the other, evaluating the statements in different possible situations. This is what we did with (1) – first we substituted 'are dogs' for 'have tails', so that the conclusion of (3) came out false while its premises were true; then we interpreted (1) in a world of Manx cats, whereby the conclusion of (1) itself came out false while its premises remained true.

Nonetheless, provided it is properly understood, the term 'representational' does significantly demarcate a semantic approach in which the interpretation of the language is fixed while the situations represented vary, different from one where one varies the interpretation within a fixed world. Moreover, the purely substitutional approach found (as Etchemendy notes, 1990: 28 ff.) in Bolzano has been extended by Tarski and later exponents to exclude limitations caused by lack of expressiveness in the language or by exigencies of the actual situation. So I will retain Etchemendy's terminology.

The crucial fact, however, is that the interpretational approach is limited – and limited in principle – in a way the representational one is not. Consider the argument:

(5) Iain is a bachelor
So Iain is unmarried.

If we substitute other expressions for 'Iain', 'unmarried' and 'bachelor', we can easily obtain an argument with true premises and false conclusion. So according to the interpretational criterion, (5) is invalid. Does it follow that (5) is really invalid? No. For there is no situation, however different from how things are, of which the premise of (5) is true and its conclusion false.

The representational account of validity says that

(R) an argument is valid if there is no possible situation where the premises are true and the conclusion false.

The interpretational account says that

(I) an argument is valid if there is no (possible) interpretation of the expressions (other than a reserved class of "logical" expressions) in the argument under which the premises come out true and the conclusion false.

So on a representational account, (5) is valid, whereas on the interpretational one the result is different; and it seems clear that the representational account is the correct one. (5) is indeed valid. The truth of the premises guarantees that of the conclusion – it is impossible for the premises to be true and the conclusion false.

We can mark the distinction here. (5) is materially, but not formally valid. Its validity depends not on any form it exhibits, but on the content of certain expressions in it. (5) is valid on account of the meanings of the expressions 'unmarried' and 'bachelor'. If we let their interpretation vary arbitrarily, we can make the premises true and the conclusion false. But that overlooks the fact that the interpretation of these two expressions is linked – their interpretation is not independent.

One might try to capture this connection formally, by insisting that the proper formalisation of (5) is something like

(6) $Fa \,\&\, Ga \vdash Ga$,

where *Fa* reads 'Iain is male' and *Ga* reads 'Iain is unmarried'. But that strategy will work only if there are logically independent semantic primitives. The famous colour exclusion problem showed that particular idea to be mistaken. Nothing red is green and vice versa, but neither is obviously more primitive than the other, and the problem recurs for any term to which each might be reduced.

Moreover, interpretational semantics cannot account for the transivity of certain relations:

(7) Iain is taller than Bill and Bill is taller than Mary
So Iain is taller than Mary.

If we are allowed to substitute other expressions for 'taller than', this valid inference will be declared invalid. But if we keep the interpretation of 'taller than' fixed then we cannot make the conclusion of

(8) No one is taller than everyone
So someone is taller than someone else,

false, and so cannot give an interpretation exhibiting true premises and false conclusion. For (8) is invalid; and its conclusion is only contingently true; we might all have been the same height. But unless we are allowed to replace some expression in it, we cannot make the conclusion false – on the interpretational account. The interpretational account, therefore, fails to capture the correct account of validity. Either we may vary the interpretation of 'taller than' freely – in which case, we give the wrong answer about (7) – or we may not vary it at all – in which case, we give the wrong answer about (8). Either way, the interpretational account gets it wrong.[2]

We may sum up our first thoughts as follows. Our reflections suggest:

(i) that not all valid arguments are valid in virtue of form, but are materially valid;

and

(ii) that validity is ultimately a matter of the impossibility of having true premises and a false conclusion, that is, ultimately a representational matter, not an interpretational one.

2. REDUCTIONISM

But this conclusion is not universally accepted. Indeed, the logical community seems roughly evenly divided – not to say, split asunder – on the question. The reason is the modality contained in (ii). What is this 'impossibility' which is referred to there? How is it to be understood – in particular, how is it to be understood if it means any more than, 'however the constituent expressions are interpreted'?

What the advocates of interpretational semantics are urging here is a reductionism about necessity. Representational semantics requires reference to various possible situations other than the actual one. If we are worried about the reality of these different possibilities – in particular, if we are worried about their epistemology, how we could know how they were constituted and what would be the case in them – we may seek to reduce their possibility to something we feel is within our grasp, namely, various different interpretations. These different *possible* interpretations can be made actual, by considering substitution-classes, and so do not suffer the remoteness of non-actual situations. We replace talk of whether a proposition might be true in a different world or situation, with how the expressions in it might be replaced by others or differently interpreted. This seems harmless enough, for there seems little to choose between, for example, conceiving of a situation in which snow is not white, but say, red, and considering the replacement of 'white' by 'red', that is, effectively conceiving an interpretation in which 'white' means not white, but say, red.

There is all the difference in the world, however. For, as Etchemendy (1988: 64) points out in his survey article on Tarski, if 'snow' comes to mean grass, 'Snow is white' will not then be true – but snow will still be white. In other words, changing the interpretation of the words changes the truth-values of the sentences, but it does not change the facts. Indeed, that is the very point of the manoeuvre, to keep the facts the same, for fear of trespassing on the unknowable, other worlds beyond our compass.

Yet that shift to interpretations loses an essential element in the analysis of validity. If Iain gets married, he ceases to be a bachelor, whereas changing the meaning of 'unmarried' has no such effect – not even that which is of course the important point here, changing the truth-value of 'Iain is a bachelor'. What is lost in interpretational semantics is the analytical linkage between expressions. For interpretational semantics properly to replace representational semantics we would need a theory which took account of these connections. But for that to be possible would require that all such connections be open to structural, or formal, articulation. The point of examples (7) and (8) was to question whether this is possible.

How can interpretational semantics – and the reductionist position on necessity – retain such an attraction if it fails so readily to provide an adequate account of validity? For it is certainly widely accepted as providing the orthodox account of the notion. There are at least two reasons. One is the numbing effect that the horror of real necessity produces – an ability to overlook the defects of reductionism if the alternative is thought to be worse. (This is a familiar aspect of any reductionism.) Another is that the counter-examples can be dismissed as not really a matter of logic: for example, the analytical connections in (5) and (7) can be seen as really a matter of meaning – it is part of the meaning of 'taller than' that is transitive, but not of logic; others, for example, the rule of infinite induction, can be set aside as again not really logic, but essentially mathematical. Logic is now seen – now redefined – as the study of *formal* consequence, those validities resulting not from the matter and content of the constituent expressions, but from the formal structure.

Etchemendy identifies a third reason for overlooking the counter-examples, an argument whose analysis will show us more about the important distinction between representation and interpretation. He calls it 'Tarski's Fallacy' (1990: chapter 6).[3] Of course, it is not a fallacy explicitly committed by Tarski. But Etchemendy does make it a plausible explanation of what lay beneath Tarski's advocacy of the interpretational approach.

The idea of the interpretational approach is this. Let us say that an argument is Tarski-valid if no variation in the interpretations of expressions other than the fixed (logical) terms makes the premises true and the conclusion (for the same interpretation) false. This much can then be deduced: it is impossible that no variation have this effect and at the same time the premises be true and the conclusion false. Tarski-validity is clearly incompatible with true premises and false conclusion. It is tempting to conclude that Tarski-validity must guarantee real validity. If Tarski-validity rules out true premises and false conclusion, does this not mean that it entails the impossibility of such a combination – that is, the impossibility of true premises and false conclusion, so that the argument really is valid?

The move is a modal fallacy, for *necessitas consequentis* (the necessity of the consequent) does not follow from *necessitas consequentiae* (the necessity of the consequence itself). Let V represent Tarski-validity and U the combination of true premises and false conclusion. Then what is true is

$$\sim \Diamond(V \,\&\, U);$$

but what is needed for Tarski-validity to ensure real validity is

$$V \rightarrow \sim \Diamond U.$$

An *ad hominem* demonstration of the illegitimacy of this inference might come from letting V stand for some contingent truth and U for its contradictory, $\sim V$. Then $V \,\&\, \sim V$ is certainly impossible, while *ex hypothesi* V is true and $\sim \Diamond \sim V$ false.

What this refutation ignores, of course, is the possibility of some analytical connection between Tarski-validity and real validity. Perhaps it is the case that the possibilities for reinterpretation are limited by the possibilities for real change, so that interpretational semantics is not completely independent of representational semantics. What we must turn to now is the general question of their relation.

3. THE SUPPRESSED PREMISE STRATEGY

In fact, Etchemendy claims to prove that, in the first order case, Tarski-validity does ensure real validity – that is, the interpretational account does not overgenerate. His argument draws on an observation of Kreisel's about informal rigour. For whereas Tarski-validity is a formal notion, precisely expressed in terms of available interpretations of expressions over an appropriate domain, real validity is an informal notion relying on an intuitive notion of impossibility.

The proof is, however, very simple. Concentrate first, says Etchemendy, on some favoured deductive system for first order logic: let $D(\alpha)$ represent '(the sequent) α is derivable in this deductive system', and let $\mathrm{Val}(\alpha)$ represent 'α is (really) valid'. Finally, let $V(\alpha)$ represent 'α is Tarski-valid'. The proof starts from the observation that the deductive system is intuitively sound – that practice has shown us that the deductive system (abstracting from its relationship to Tarski-validity) does not overgenerate:

$$(\forall\alpha)(D\alpha \rightarrow \mathrm{Val}\alpha).$$

We now appeal to a formal result, the completeness of first order logic (with respect to interpretational semantics):

$$(\forall\alpha)(V\alpha \rightarrow D\alpha).$$

We immediately obtain, by transivity,

$$(\forall\alpha)(V\alpha \rightarrow \mathrm{Val}\alpha),$$

assuring us that Tarski-validity (V) does not exceed real validity (Val).

This is a quite extraordinary argument. What is it supposed to show? Etchemendy has led us to consider Tarski's account of validity, and has given reasons for us to be sceptical about its extensional correctness – *via* attacks on its intensional formulation. He then invites us to take some deductive consequence relation, provably equivalent to that suspect account of validity, and concludes from its intuitive soundness that it really is sound.

Nonetheless, the following tempting objection is mistaken. Etchemendy's proof purports to show that every argument valid in first order (classical) logic is really valid. But there are many counter-examples to this claim – or at least, even if they are not accepted as counter-examples, Etchemendy's argument is far too quick with what is a matter of deep contention. Take, for example, the principle *ex falso quodlibet*, from A and $\sim A$ to B, or the principle of double negation, from $\sim\sim A$ to A. Surely the best that Etchemendy's argument can do is to focus our attention back on his premise, namely, that D is intuitively sound – that $D\alpha$ entails Valα. The proof seems to add nothing to that claim.

But this is not so. What Etchemendy's (and Kreisel's) point shows is that Tarski-validity is extensionally safe for any provably complete deductive system which one believes is intuitively sound. That is, whatever your scruples, let $D(\alpha)$ represent 'α is provable in my preferred first order logic' – whether classical, intuitionistic, relevant or whatever. Then, if you have a completeness proof for this logic relative to its Tarski semantics, its intuitive soundness (for you) will carry over to its Tarski soundness, that is, Tarski-validity, suitably defined – i.e. the use of interpretational semantics – will not proclaim any really invalid arguments to be valid. Tarski-validity will be safe in the sense that it will not overgenerate.

It is not, however, a result with which to get carried away. It assumes we already have a Tarski-complete proof procedure – so that demonstrating validity is not a particular problem. The purpose of the semantics is to show invalidity. But what guarantee have we that Tarski-invalidity entails real invalidity? We saw in §1 that there is no guarantee at all, that it does not.

Nonetheless, Etchemendy sets out to adapt Kreisel's remark to show that, again for first order logic, we do not have undergeneration either – that real validity entails Tarski-validity. However, he himself points out that his result has limited significance, applying only to the common logical truths of a class of languages, that is, where we have already abstracted from logical truth particular to a certain language, in other words, where the logical truth in question does not depend on the particular meanings of any of the terms involved which are not common to all the languages.

We saw in §1 that this overlooks vital features of validity. Arguments (5) and (7) are valid, yet their validity depends on the particular interpretation of 'married', 'bachelor' and 'taller than'. Allowing their interpretation to vary arbitrarily loses sight of this dependency; fixing it prohibits consideration of other ways in which validity can fail. Interpretational semantics cannot do justice to the possibilities of representation and the facts of validity. An argument is valid if and only if there is no possible situation where the premises are true and the conclusion false.

The reductionist, or formalist, has another strategy open to him, however. Somehow this obscurantist reference to real possibilities must be removed. Perhaps the way to do it is to demand that these analytical connections which defy formalisation and undermine interpretational semantics be explicitly stated. That is, if arguments such as (5) and (7) are valid, they are only really valid when the suppressed premise, revealing the underlying connections, is made explicit. In the case of (5), what is needed is an extra premise stating that all bachelors are unmarried; in the case of (7), a premise to the effect that 'taller than' is transitive.

The suppressed premise strategy appears in Etchemendy (1990, p. 68) in the guise of 'cross-term restrictions'. This is a somewhat misleading name in general. In the case of (5), we require that the terms 'bachelor' and 'unmarried' be linked, so that the interpretation of 'bachelor' be restricted to a subset of that of 'unmarried'. That is clearly a 'cross-term' restriction – the restriction on the interpretation of 'bachelor' is relative to that of 'unmarried'. However, the restriction needed on 'taller than' in (7) (that it be interpreted only by transitive relations) is 'cross-term' only in a degenerate sense – it simply and absolutely restricts the interpretation of 'taller than' with no reference to any other term.

There is something very puzzling about the suppressed premise strategy, however. What exactly is added to the argument by making the hidden premise explicit? It may well have psychological value, in clarifying the reason for the argument's validity. But it cannot turn an unsound argument into a sound one.

For suppose (5) were in fact invalid. Then, at least for Iain, that he is a bachelor would not entail that he was unmarried (so not all bachelors would be unmarried). The added premise would, therefore, be false. Hence the expanded argument:

(9) Iain is a bachelor
All bachelors are unmarried
So Iain is unmarried,

would have a false premise.

Of course, the point of the suppressed premise move is that the added premise should be true, as in the case of (5) it is. Indeed, it must be logically true. This is not brought out clearly by (5), for if its added premise, 'All bachelors are unmarried', were not logically true it would be false. Contrast:

(10) All bachelors are unmarried
So Iain is unmarried.

If we add the extra premise that Iain is a bachelor, we obtain a valid argument. But (10) is not valid, for the premise can be true and the conclusion false. So it does not suffice simply to add true premises to an argument to show that it is really valid. The added premises must be logically true.

This may recall the notion of enthymeme, which is used in this broad way to refer to any argument, such as (10), which can be converted to a valid argument (9) by the addition of a true premise. But that is not the current strategy, which is to add premises which are logically necessary. Our focus is on arguments whose premises cannot be true and conclusion false, and yet which do not instantiate a valid form. For them, the added premise will be logically true. Another worry arises, nevertheless – or rather, the original worry in a new guise. For the modality which the reductionist sought to remove by the suppressed premise strategy has now shifted to the added premise. Which added premises are acceptable and which not – which are the logical truths?

The interpretational approach is often coupled with the idea that logic is topic-neutral – what is purely logical abstracts away from the specific content and presents a formal schema. But will the supposed divide between the specific and the general – between content and form – support the weight put on it?

It will not. Interpretational semantics has to mark off a class of expressions (logical constants) which are immune to variation. For if all expressions could vary their interpretation, only the most trivial examples would be accounted valid. But examples (5) and (7) above show that the class cannot in fact bail out the interpretational approach.

In fact, one can see dissatisfaction with this problem lying behind Quine's attack on the notion of analyticity. Hart (1991) recalls the debate between Tarski and Quine (believers in 'plain truth') and Carnap (a believer, as Hart puts it, in 'fancy') over the nature of logical consequence. As long as one recognises analytical, and so logical, connections between expressions, interpretational semantics will be guaranteed to undergenerate. The only solution available to an advocate of 'plain truth' is to disparage such connections, suggesting they are open to revision, and so not logically compelling after all.

Yet this strategy must fail – on pain of dismissing logic altogether. For logic requires a separation between simple truth – whether the constituent statements of an argument are true or false – and modality – whether the (possible) truth of the premises guarantees that of the conclusion. The reductionist strategy transfers this distinction from the statements themselves (are the connections analytic or synthetic?) to that of form and content (separating the logical constants – whose interpretation may not vary – from the substantive content, which is removed). But if we refuse to recognise a set of inferential connections between substantive expressions as analytical, neither can we recognise a set of expressions as logical constants, as resistant to reinterpretation with everything else.

'Fancy,' or logical necessity, lies at the heart of logical consequence. Valid arguments are those whose premises cannot, logically, be true when their conclusions are false. This requires that we remain prepared to study the inferential connections in (5) and (7), as much as those in (1) and (3).

4. REPRESENTATIONS TRIUMPHANT

To return to the suppressed premise strategy, however: the puzzle remains. The extra premise is strictly redundant. For if the original argument were invalid, the added premise would not be logically true. Given that it is logically true, it follows that the unexpanded argument was already valid. Hence it was (logically) unnecessary to add the extra premise.

The point is reminiscent of a familiar puzzle about *Modus Ponens*, that the major premise is either false or unnecessary:

$$A, \text{ if } A \text{ then } B/\text{so } B.$$

If the major premise is true, then *B* follows from *A*, and so the major premise is redundant. So either the major premise is false, or not needed at all. Either way, it contributes nothing to the argument.

Of course, what we are looking at in the suppressed premise strategy is not whether *B* simply follows (perhaps contingently on other assumptions) from *A*, but whether the conclusion is entailed by the premises – whether the argument from *A* to *B* is valid. So our problem concerns the logical, and not merely contingent, truth of the major premise. Nonetheless, the problems have a striking structural similarity. If the added premise is available (i.e. is logically true) then it is not needed (the argument is valid without it); so if the argument is invalid without it, the additional premise is not logically true, and so not available to be added.

There is another, and more instructive, way of seeing the point about *Modus Ponens*, however. A conditional records the fact that we may correctly proceed from its antecedent to its consequent. That is all there is to the meaning of conditionals. We might call it a redundancy theory of conditionals, on analogy with the redundancy theory of truth. Not that the correctness of this view of conditionals depends on that for truth. Indeed, as we saw in §2, the truth-predicate is not always redundant, as in 'If "white" comes to mean red, "Snow is white" will not be true', where '"Snow is white" will not be true' cannot be replaced without loss of truth by 'snow will not be white'. Nonetheless, the idea is the same: conditionals do not make reference to some strange species of conditional fact, but simply record successful passage from one fact (the antecedent) to another. Indeed, just as the redundancy theory of truth does not claim that all talk of truth can be eliminated, neither does the redundancy theory of conditionals. The redundancy theory of truth claims that '"All cats are animals" is true' says no more and no less than 'All cats are animals', and here talk of truth is eliminable; but in 'What Iain said is true' it is not. So too with conditionals.

This is the burden of normalisation. Any proof in which an application of Conditional Proof follows one of *Modus Ponens* (with the same formula as conclusion of one and major premise of the other) is more complex than is necessary. A reduction step is possible:

$$\begin{array}{c}\dfrac{\begin{array}{c}\Pi_1\\ A\end{array}\quad\dfrac{\begin{array}{c}[A]\\ \Pi_2\\ B\end{array}}{\text{If } A \text{ then } B}}{B}\\ \Pi_3\end{array} \Rightarrow \text{(reduces to)} \quad \begin{array}{c}\Pi_1\\ A\\ \Pi_2\\ B\\ \Pi_3\end{array} \qquad (1)$$

Here the statement of 'If *A* then *B*' is redundant. Π_2 not only shows that B is derivable from *A*, but actually derives it.

Where conditionals come in useful is where we don't actually have Π_2, the crucial derivation, only a report of its existence. The case parallels that for truth. Truth enables us to carry various reports around under certain descriptions (e.g. 'what Iain said') without all the bothersome detail. Similarly, conditionals enable us to transmit a record of proof without its detail. If *B* does indeed follow from *A*, then in some sense we need neither the proof, nor the record, validly to move from *A* to *B*. But in order to be *assured* that the step is valid, we can rely either on the proof, or, if we simply have a reliable report – 'if *A* then *B*' – that the step is valid, we can use *Modus Ponens* to articulate it.

The *Modus Ponens* puzzle is dramatic. But the puzzle also affects the other connectives. The introduction rule for a connective states the ground on which statements containing that connective can be made. The Cut principle (or normalisation) then links the introduction rule to an elimination rule in such a way as to guarantee conservativeness – that is, adding a connective via its introduction rules does not affect the provability of statements lacking that connective. So *&E*, $\vee E$ and so on are equally redundant – when tied to the grounds for assertion of the major premise a reduction procedure eliminates the inference and the major premise. But this redundancy is theoretical, not real. The logical connectives are useful in bundling information – that *B* follows from *A* in the case of conditionals, that one of *A* or *B* is true in the case of disjunction, and so on. We may no longer have access to the derivation of *B* from *A*, or to the information about which of *A* or *B* is true, respectively, but 'If *A* then *B*' and '*A* or *B*' carry that assurance forward so that if in the first case, we come to find out that *A* is true, we can infer *B*, or if in the second, we should find out that *C* follows from *A* and from *B* equally, that *C*. The logical connectives import no information of their own, but serve to record combinations of other facts.

Does this response to the puzzle about *Modus Ponens* help us with the puzzle about the suppressed premise strategy? It tells us two things: first, that the added premise adds psychological perspicuity, in spelling out why the argument is valid. In that sense, it is not redundant, any more than the conditional is in *Modus Ponens*. But at the same time, it is not essential. It would be wrong to say that an argument is valid in an extended sense if an expanded argument, containing an extra, logically true premise, is really valid, if that should imply that arguments valid in the extended sense are not really valid. If the major premise of an application of *Modus Ponens* is true, then its conclusion does indeed follow from its minor premise, and the major premise is redundant – except for assuring us of that consequence. Similarly, if the extra premise added to an argument is logically true, then the conclusion does follow from the remaining premises – so arguments valid with added premises really are valid. What the extra premise does is to assure us of – and perhaps explain to us – the validity of the unexpanded argument.

Our puzzle about the suppressed premise strategy should now be solved. The extra premise adds clarification, but it does not serve to turn an invalid argument into a valid one. The strategy cannot, therefore, be used as a defence of interpretational semantics. Recall that the problem we are faced with is that interpretational semantics cannot be controlled so as to match, and serve as a replacement for, representational semantics. Allowing certain expressions to vary their interpretation undergenerates; fixing their interpretation overgenerates. The suppressed premise strategy tries to close this gap by restricting the variation so that it varies no further than the validity of the argument (if it is valid) may permit.

In this it may succeed – indeed, as stated, it must succeed. For if the argument is valid, then there must be some fact about the interpretation of premises and conclusion which rules out the possibility that the premises are true and the conclusion false. Explicating the fact which rules out the possibility of true premises and false conclusion will result in a logically true statement whose truth will serve to restrict the permissible interpretations of premise and conclusion, so that the truth of the conclusion cannot drift away from that of the other premises without falsifying the premise which has been added. In other words, representational semantics guarantees that the suppressed premise (or "cross-term") extension of interpretational semantics will succeed.[4]

5. LOGIC

What the believer in the suppressed premise approach to material validity must convince us of, is that the restriction on possible interpretations can be formulated and stated as a separate (in itself insufficient, but jointly sufficient) logically true added premise. It may be that, provided there is no limit on the richness of the resources available, this can always be done. The important point, however, is that the added premise does not turn an invalid argument into a valid one. It turns a materially valid argument into a formally valid one, by suitably restricting the range of formal counter-examples. For validity is at root a matter of what can be represented. It is representability which makes the (original) argument valid, and which makes the added premise logically true. There are arguments which are valid and not formally valid. To show invalidity we must, in the end, turn to representations in which the premises are true and the conclusion false.

Logic is the study of valid inference. Sometimes that validity is purely a matter of form – elaborating the way the logical connectives bundle information together, and drawing conclusions from it. But that does not exhaust the study of validity.

Commensurate with this wider brief for the logician, taking him out of the closet of the study of pure form into the wider world of material consequence, is a serious limitation on what can be expected of the logician. What is logic for? Logic can tell us what counts as a valid argument, what the criterion for validity is – what is needed for an argument to be valid. But only rarely can it tell us which arguments actually are valid. Indeed, where it can tell us this, is precisely in the study of pure form. Beyond that, we are faced with a need to discover whether the premises can be true and the conclusion false. To decide this may be a complex issue in metaphysics, for example, or in mathematics. Since 2^{127} is even, we may ask whether it follows that it is the sum of two primes. If Goldbach's Conjecture is true, it does. But the logician cannot answer that question. The logician's task is to observe that Goldbach's Conjecture will serve to entail the conclusion – for if every even number[5] is, as a matter of mathematical necessity, the sum of two primes, then it is impossible that 2^{127} be even and that 2^{127} not be the sum of two primes. Similarly, if determinism is incompatible with free will, one may validly infer the falsity of determinism from the presence of free will (or contrariwise). This may seem to invite the metaphysician dangerously into the logical arena. Should logic not be topic-neutral, and insulated from such intrusion? But that is not so. The topic-neutrality of logic need not mean there is a pure subject matter for logic; rather, that the logician may need to go everywhere, into mathematics and even into metaphysics. Nor need we deny that making the suppressed premises explicit is not a useful task. Where possible, exhibiting a pure form of argument can yield clarification, and can show validity by abstracting a valid form.

What must be acknowledged is that belief that every valid argument is valid in virtue of form is a myth, and exclusive concentration on the study of pure forms of argument does a disservice both to logic and to those who can be helped by it. Validity is a question of the impossibility of true premises and false conclusion for whatever reason, and there are arguments which are materially valid and where that reason is not purely logical.

NOTES

[1] Indeed, Etchemendy (1983: 326 ff.) called it the "substitutional/interpretational theory".

[2] One strategy for the formalist here is to reduce the preordering on people to a partial ordering on heights, so reducing the transivity of 'taller than' to that of 'greater than'. But the same problem now arises for 'greater than': if it is not a logical term, interpretational semantics will allow its interpretation to vary so that the analogue of (7) is deemed invalid; while if its interpretation is kept fixed, the analogue of (8) comes out valid – but it is not, for everything might have been the same size.

[3] In his (1983: 330) he called it 'Bolzano's Fallacy'.

[4] Does adoption of the representational account of validity, asks the referee, commit me to the belief that a necessarily true conclusion follows from any set of premises whatever? For if an argument is valid if and only if it is impossible for the premises to be true and conclusion false, then surely if the conclusion cannot be false, it cannot be false in conjunction with any set of premises, and so follows validly from them. I have deliberately kept the body of this paper neutral, as I see it, between the classicist, the intuitionist, the relevantist and others. I do not, however, believe that the modal – that is, the representational – account of validity has this consequence. I have discussed the matter at some length elsewhere (Read 1988: especially chapter 3). In brief, Tarski's generic conditions on consequence, Weakening (or Monotonicity) and Cut (with its special case, Suppression) are correct only if properly controlled. That control requires that premises can be augmented only in a weak way, by a loose kind of extensional binding; but necessary truths can be suppressed only when bound in a strong way, by a tight and intensional binding. So the tempting move to irrelevance turns on an equivocation and lack of care over the interpretation of conjunction in the formula "impossible for the premises to be true and conclusion false". See also Read 1981.

[5] Greater than 2.

REFERENCES

Etchemendy, J. 1983: "The Doctrine of Logic as Form", *Linguistics and Philosophy* **6**, pp. 319–34.

Etchemendy, J. 1988: "Tarski on Truth and Logical Consequence", *Journal of Symbolic Logic* **53**, pp. 51–79.

Etchemendy, J. 1990: *The Concept of Logical Consequence*, Cambridge, Mass.: Harvard University Press.

Hart, W. D. 1991: "Critical notice of Etchemendy (1990)", *The Philosophical Quarterly* **41**, pp. 500–5.

Read, S. L. 1981: "Validity and the intensional sense of 'and' ", *Australiasian Journal of Philosophy* **59**, pp. 301–7.

Read, S. L. 1988: *Relevant Logic*, Oxford: Blackwell.

Department of Logic and Metaphysics,
University of St. Andrews,
St. Andrews, Fife KY16 9AL
Scotland