CANADIAN JOURNAL OF PHILOSOPHY
Volume 25, Number 2, June 1995, pp. 283 - 292

Etchemendy and Logical Consequence

GRAHAM PRIEST
University of Queensland
Brisbane, Queensland
Australia 4072

I Introduction

Logical consequence is a notion that every person who reasons must possess, at least implicitly. To give a precise and accurate characterization of this notion is the fundamental task of logic. In a similar way, the notion of effectivity is a concept that anyone with a basic training in mathematics possesses, and the most fundamental task of a theory of computability is to give a precise characterization of this notion. The problem concerning effectivity was solved (at least to the satisfaction of most people) in the 1930s, almost as soon as it was raised, by the work of Turing, Church, and others. By contrast, the correct and precise characterization of logical consequence has been hotly contested through the two and a half thousand-year history of logic.

Despite this, for the last fifty years or so there has been relative consensus on the issue — at least as far a deductive validity goes. It is true that this has been challenged by heretics of intuitionist and paraconsistent variety; but the power of the first year logic text-book is a very hard thing to take on. The recent book by Etchemendy attempts this difficult task.[1] In this note I want to discuss Etchemendy's views and related issues.

---

1 J. Etchemendy, *The Concept of Logical Consequence* (Harvard, MA: Harvard University Press 1990). References are to this book unless otherwise indicated.

II The Tarskian Orthodoxy

A first stab at a definition of deductive validity might go as follows. An inference is valid iff it is impossible that the conclusion is not true, given that the premises are true. The problem about framing a rigorous account of validity is the problem of how to make the notions involved in this homely truth precise. A first step is to interpret the 'given that' as a simple conditional: if all the premises are true then, necessarily, so is the conclusion. This may seem a straightforward step, but it already takes us down the path from a simple truism to a substantial theory. To see this, just compare it with the corresponding situation for inductive validity. The corresponding first stab at a definition for inductive validity would be: it is *improbable* that the conclusion is not true, given that the premises are true. The natural way to cash out the 'given that' here (whether or not this is correct) is not as a conditional, but as a conditional probability; so the definition becomes: an inference is inductively valid when the probability of the conclusion, conditional upon the premises, is high. The fact that orthodoxy cashes out the 'given that's in such different ways in apparently parallel cases should at least give one pause to wonder why.

At any rate, once we have sorted out how to understand the 'given that' in the definition of deductive validity, the next issue is how to cash out the notion of necessity involved. At this point the orthodox account makes a bold move. Necessity is to be cashed out as quantification over all things of a certain kind. This approach can be found *in nuce* in a number of earlier writers, but its site of crystallization is Tarski's famous 1936 paper.[2] Essentially, it goes as follows.

Let φ be any formula of some well-defined language. It is assumed that we can divide all the symbols of this language into distinct grammatical categories. Fix on some set of those symbols and call them *logical constants*. Replace every non-logical-constant in φ with a new symbol, to be called a *parameter*, symbols of different grammatical categories being replaced by different styles of parameter. Call the result a *form*. An *assignment* assigns to every parameter an entity appropriate to its category (an object, set, etc.). An assignment *satisfies* a form iff the sentence obtained by substituting, for each parameter, a word naming or otherwise expressing the entity assigned to that parameter, is true. Finally, an inference is valid iff: if v is any assignment that satisfies all the premises,

2 'On the Concept of Logical Consequence.' The paper was reprinted in English in 1956, as ch. 16 of *Logic, Semantics and Mathematics* (Oxford: Oxford University Press 1956).

v satisfies the conclusion. A sentence is a logical truth iff every assignment satisfies it. Note that the validity of the inference (or the logical truth of a formula) is expressed as a certain universal generation over all assignments.

III Preliminary Comments

Tarski's definition of validity does not deliver the text-book notion of validity immediately, for several reasons. The first is that to apply it we have to agree on *which* notion of conditionality is being used in the definition. The standard assumption (and one not questioned by Etchemendy) is that it is the material conditional; but this is not at all obvious, since it leads directly to the highly counter-intuitive result that if something is logically impossible, everything follows validly. Arguably, it is much more natural to take the conditional to be non-material, as do some non-classical logicians.[3]

The second reason is that even given that this is settled, the notion of validity piggy-backs on a notion of satisfaction. The notion itself is unproblematic, but the question of what satisfies what, is not. We can all agree, for example, that:

> Thatcher satisfies "x is a prime minister of Britain" iff Thatcher is a prime minister of Britain

and that:

> Thatcher satisfies "if x is a prime minister of Britain then Brisbane is in Queensland" iff (if Thatcher is a prime minister of Britain then Brisbane is in Queensland).

And we will all agree that Thatcher *is* a prime minister of Britain. But we may not agree on the truth of the conditional: if Thatcher is a prime minister of Britain then Brisbane is in Queensland. And a similar point might be made with respect to other constructions, e.g., ones using 'necessarily,' 'probably,' or even 'or.'

Tarski, of course, gave a recursive definition of satisfaction for languages of certain kinds; and it might be thought that this overcomes any uncertainty. It does not; but merely shifts the problem into the metalan-

---

3 E.g. S. Read, *Relevant Logic* (Oxford: Blackwell 1988); G. Priest, 'Boolean Negation and All That,' *Journal of Philosophical Logic* **19** (1990) 201-15, 207f.

guage. The Tarskian homophonic truth conditions for the conditional are: assignment v satisfies $\langle\alpha\rightarrow\beta\rangle$ iff (if v satisfies $\langle\alpha\rangle$ then v satisfies $\langle\beta\rangle$). Thus any uncertainty about the truth of an object level conditional will transfer to the metalinguistic conditional. Nor is there anything intrinsic to the recursive definition that requires this conditional to be material, as a number of people have noted.[4]

The text book account of validity resolves the issues above in familiar ways, by enforcing a certain understanding of the conditional and other logical particles. The adequacy of these ways (at least as far as the conditional goes) is highly contentious (as the references cited show). However, Etchemendy passes over these issues in silence and so I shall not discuss them further.

IV Cross-term Restrictions

Even with the above issues resolved in the orthodox way, Tarski's account still does not appear to give the text-book account of validity, as Etchemendy notes (ch. 5). The problem concerns quantifiers. The standard model-theoretic semantics treat quantifiers as logical constants; yet it is crucial in defining validity that the quantifiers be allowed to range over different domains. Hence the quantifiers can not be treated as simple logical constants. The obvious solution is that the standard quantifiers, $\supset x$ and $\exists x$, are tacitly being thought of as of the form $\sqsupseteq x\varepsilon D$ and $\exists x\varepsilon D$, respectively, where D is a parameter.

The problem with this, as Etchemendy points out, is that it invalidates standard logical inferences. Universal Instantiation, for example, is now really of the form: $(\supset x\varepsilon D)\varphi(x) \vdash \varphi(a)$. And this is invalid since a may not be in D. To get around this problem, we have to insist that there be certain restrictions on what objects are allowed to be candidates for satisfaction. For example, an assignment must be allowed to assign to 'a' only a member of the set which it assigns to 'D.' Etchemendy calls these 'cross term restrictions,' and rightly points out that these are an unmotivated departure from the original definition, and hence a tacit admission of its inadequacy.

It seems to me, however, that the problem reflects ill not so much on the definition of validity as on the ways standardly chosen to abstract natural language into a formal language. In English, a quantifier is

4 For example, it might be some relevant conditional, as in G. Priest and J. Crosthwaite, 'Relevance, Truth and Meaning,' in J. Norman and R. Sylvan, eds., *Directions in Relevant Logic* (Dordrecht: Kluwer 1989).

always followed by a count noun: every dog, some man, etc. Moreover, because of this, Universal Instantiation should be expected to fail. <Every dog is four-legged; hence Napoleon is four-legged> is obviously invalid: for the conclusion to follow we need the information that Napoleon is a dog. A more adequate formalization of this would use a formal language in which every quantifier is bounded in the obvious way by a domain parameter, D. We might write quantifiers thus: ⊃Dx, ∃Dx. In the semantics, the parameter would be assigned some set of objects, and the bounded quantifier would range over this set.

One might object to this procedure that at least the inference <Everything is a dog; hence Napoleon is a dog> is logically valid. But this intuition can be accommodated by taking 'everything' to be 'every thing,' where 'thing' is a logical constant that is assigned the set of all objects. Etchemendy, in fact, considers this, and rejects it on the ground that it makes ∃Tx∃Tyx≠y a logical truth (where 'T' is the logical constant *thing*). This raises the issue of overgeneration, which I will take up later, so let us leave the discussion here for the present.

V Epistemic Considerations

The problem about cross-term restrictions is only a minor problem for the Tarskian account, according to Etchemendy. The most telling form of criticism for any definition is the production of a counter-example. It is such criticism that takes center stage in Etchemendy's book. But before we turn to this, I want to discuss some further considerations adduced by Etchemendy which may be thought of as an argument against the Tarskian account (89f).[5] These are as follows.

If the validity of an inference is to be identified with the truth of a universal generalization, then we cannot know that an inference is valid unless we know this generalization to be true. But we cannot know that this generalization is true unless we know that its instances are true; and we cannot know this unless we know that every instance of an argument form is materially truth preserving. Hence, we could never use the fact that an argument form is valid to demonstrate that an argument *is* materially truth preserving. Thus the prime function of having a valid argument would be undercut.

---

5 They are certainly used this way by some writers, e.g., G. Priest, 'Two Dogmas of Quineanism,' *Philosophical Quarterly* **29** (1979) 289-301.

This argument, in fact, fails. This is simply because epistemic order need not coincide with definitional order. To see this, just return to the definition of effectivity provided by Church and Turing, and note that we can be sure that a process is algorithmic before we check to see whether we can produce a Turing machine to perform it. If we could not, the Church/Turing thesis would have no bite. Similarly, even if the definition of validity is correct, we may be able to determine that an inference is valid without consulting the definition. This, of course, raises the question of how we tell that an inference *is* valid (or that a procedure is algorithmic). But that is another issue.

VI Undergeneration

Let us now turn to Etchemendy's central objections to the orthodox account. According to Etchemendy, the orthodox account will, in general, both make intuitively valid inferences invalid (undergeneration) and make intuitively invalid inferences valid (overgeneration). And where, on any particular occasion, it does get the extension of the notion of validity right, it will do so only by happenstance.

Undergeneration will arise when the validity of an inference is due to the meaning of a word that is made a parameter. Thus, for example, if we take predicates and names to be non-logical-constants, the inference <Leslie is a brother; hence Leslie is a male> is invalid. And if we try to avoid this conclusion by taking predicates to be logical constants, the result is overgeneration. For the inference <Leslie was a US president; hence Leslie was a man> then becomes valid.

The problem here is one of ancient ancestry. The ancient response to it (which Etchemendy does not consider) is to grit one's teeth and insist that the original inference is just invalid. What *is* valid is the inference with the suppressed premise: all brothers are male. The original inference is an *enthymeme*. The response is, to a certain extent, counter-intuitive: the original inference did seem valid as it stood. However, there is no reason why a theoretical account of validity should show all our original intuitions about it to be correct, any more than a theoretical account of motion should show all our original intuitions about what is in motion to be correct. And at least this account explains our mistaken intuitions: we confuse the argument with a similar valid one.

VII The Omega Rule

Etchemendy often uses another example of undergeneration (e.g., p 100). This is the infinitary inference usually called the omega rule:

> 0 is P, 1 is P, 2 is P,... Hence all numbers are P.

Intuitively, he says, this is a valid inference; and so did Tarski. Yet if we do not take numerals and 'number' to be logical constants, the inference is clearly invalid under the orthodox definition. An obvious suggestion as to how to get around this, is to insist that these symbols be taken as logical constants. This raises the issue of what, exactly, constitutes a logical constant. However, independently of that issue, the suggestion is problematic. For the sentence 'there are two distinct numbers' would then contain only logical constants, and hence (vacuously), be satisfied by all assignments. Thus it would be a logical truth, which seems counter-intuitive.

Again, however, the problem can be avoided simply by denying the validity of the omega rule. As before, we can take it that what is valid is the closely related enthymeme with suppressed premise: if C then all numbers are P; where C is the conjunction of: 0 is P, 1 is P, etc. Of course, this is an infinitary sentence, but the rule was an infinitary rule; hence anyone who has finitist objections to the solution should have had finitist objections to the problem.

VIII Overgeneration

So far, we have discussed only the problem of undergeneration. We must now face the tougher problem of overgeneration. This is supposed to arise when the universal generation which states the validity of an argument, though true, is so only 'accidentally.' An example which Etchemendy takes to show the problem (e.g., 123f.) is the following.

Consider the language of second order logic, where the logical constants are taken to be the usual ones. It is well known that there is a formula, φ, of this language, containing only logical constants and variables, that expresses the Continuum Hypothesis (the claim that any uncountable set of real numbers has the same size as the set of real numbers). Moreover, it is also well known that this sentence is categorical (has only one model) up to isomorphism. What this means is that according to the Tarskian definition φ is logically true iff the Continuum Hypothesis is true. So either it or its negation is a logical truth. But, the argument continues, the Continuum Hypothesis is a substantive claim about a particular domain, the domain of sets. Hence neither φ nor ¬φ should come out as a logical truth.

Of course, this raises the hoary old issue about whether set-theory is logic. But the argument can be seen to fail quite independently of this issue. Recall that in section 3 we saw that for the Tarskian account of validity to be applied correctly, all quantifiers should be bounded by a domain parameter. Now if one takes φ and bounds all first order quantifiers by a domain parameter, and all second order quantifiers by another domain parameter, the resulting sentence is not categorical, and so the argument fails.

The reason why is, of course, that in the standard interpretation of second order logic the domain of the second order quantifiers is not arbitrary, but is required to be the power set of the domain of the first order quantifiers. If we want to obtain this effect now that quantifiers have been treated properly, we must enforce it. There are numerous ways of doing this, but they all come to the same thing in the end. We introduce a new second-order functor P (power set) and require that if the first order quantifiers are bounded by D, the second order quantifiers are bounded by P(D). (In English, we would simply say 'every *set of* Ds' for ⊃P(D)X, etc.)

If we rewrite φ in this way, what is the logical status of the resulting sentence? This depends on whether we take 'P' (or *set of*) to be a logical constant. If, as tradition would have it, we do not, then it can be interpreted as any second order functor (e.g., *set of finite subsets of*). The required linkage between first and second order quantifiers is therefore broken, and the sentence is not logically true or logically false. If, on the other hand, we treat P as a logical constant for some reason or other, the sentence is logically true or logically false. But this is hardly surprising: if we treat 'power set' or 'set of' as fixed in meaning, then of course certain claims about sets are liable to turn out to be logical truths.

IX Identity

Another example that Etchemendy uses to illustrate the situation (e.g., 111f.) is more difficult. Take the logical constants to be the usual ones, and consider the claim that there are at most n cats, where this is shorthand for the usual sentence of first order language starting with n+1 universal quantifiers. If the cardinality of the cosmos is n or less, this will turn out to be a logical truth, which seems wrong. One might reply that the cardinality of the cosmos, all things (including the objects of mathematics) considered, is greater than n; but this would be to miss the point. The point, as with the previous example, is not that the sentence might be a logical truth, but that its being a logical truth ought not to depend on some substantive fact about the world: physical, mathematical or other.

This argument obviously puts a lot of weight on the notion of a substantive fact; and a major problem is that it is not very clear what this is. That there are so many physical objects in the cosmos is clearly such a fact. That there are so many existent objects (where this might include mathematical objects) is less clearly so (at least to anyone who takes mathematics to be part of logic). But that there are so many objects *in toto* does not appear to be a substantive fact at all. At least arguably, this collection includes not only all actual objects, but also all possible, and maybe even impossible, objects. That there are exactly so many of these hardly seems to be a matter of any contingency or variability. Nor is it a truth about some special domain or other in any obviously objectionable sense.

Of course, it could be replied that possible objects are not objects at all. And it must be admitted that, in general, possible Xs are not Xs (a potential killer is not a killer; a possible accident is not an accident). But a possible object is something that can be thought about, quantified over distinguished from other things, attributed properties (if only the property of being a possible object); and if all these do not ensure that it is an object, I do not know what could.

It could, I suppose, be argued that although all objects exist, many possible objects do not exist (the merely possible ones); and hence that there are more possible objects than objects. The claim that all objects exist is one that was made familiar by Quine in the form of the maxim 'to be is to be the value of a bound variable.' Whether or not this claim can be substantiated (I do not believe so), it suffices here to note that someone who claims this cannot also claim that there are possible objects that do not exist. Hence this argument will not work.

Conceivably, there are others, but let us not pursue this here. Provided we take the domain of quantification to be *all* objects — actual and possible — then it is not at all clear that the Tarskian account delivers the wrong answers about validity. In particular, every sentence of the form: there are n things, *is* a logical truth. (See the end of section IV.) Nor is it clear that it delivers the right answers for the wrong reasons. The modality involved seems just what is required if Tarski's reduction of a modal notion to a quantifier is to be right.

X Conclusion

Etchemendy gives a few other examples of overgeneration, but they can all be handled by the replies developed in the last couple of sections, or combinations thereof. Assuming that no more counter-examples can be found (which is not at all obvious), this shows only that the orthodox definition has the right extension. It could still be objected that the definition is wrong because it has the wrong intention. I doubt that such an objection can be made to stick, however. The definition of an algorithm in terms of Turing machines hardly gives the meaning of the preexisting word 'algorithm.' But this is not seen as an objection as long as we consider ourselves as giving a real definition, as opposed to a nominal definition. The same is true of the definition of validity.

For this reason, Etchemendy's case must rest largely on the persuasiveness of his counter-examples. I have argued that these are not, in the end, as telling as he claims; though, as I have also indicated, there is room for dispute. The dispute will involve substantial metaphysical issues. Much as one might regret this (if one had hoped that the foundations of logic could be relatively independent of such murky areas), it is inevitable. This explains, incidentally, why logic has never divorced itself from philosophy to become a part of another subject (such as mathematics) or an autonomous discipline, as so many other areas have. Nor, for the same reason, will it ever do so.

At any rate Etchemendy has provided a much needed reminder that the orthodox account of validity (and probably any account) is problematic. Its contentious assumptions go well beyond those that Etchemendy has located, as I have indicated, though Etchemendy has taken on a very central assumption that even those working in non-classical logics have pretty much taken for granted. It is to be hoped that Etchemendy's book will stimulate discussion of all the contentious assumptions packed into the orthodox account of validity. Too many have taken these for granted for too long.

Received: October, 1993