Logical Pluralism

JC Beall
School of Philosophy
The University of Tasmania

Greg Restall
Department of Philosophy
Macquarie University

Version of September 15, 1998

"... in considerations of a general theoretical nature the proper concept of consequence must be placed in the foreground."

Alfred Tarski [58, page 413]

Abstract: A widespread assumption in contemporary philosophy of logic is that there is *one true logic*; that there is an unequivocal answer to any question of the deductive validity of any determinate argument. We reject this account of logic, and in this paper we propose an alternate view, which we call *logical pluralism*. According to logical pluralism, there is not *one* true logic, there are *many*. There is not always a single answer to the question "is this argument valid?"

This paper motivates logical pluralism by way of an analysis of disagreement within the tradition of classical first-order logic, and between this tradition and two of its rivals, relevant and intuitionistic logic. We argue that debates between these "rival" logics — and between different accounts of the one logic — motivate a pluralist account of logical consequence.

We end the paper by defending logical pluralism against a number of different objections to the view.

1 Logic, Logics and Consequence

Anyone with a nodding acquaintance with contemporary Logic knows that there are many so-called logics.[1] But are these logics rightly so-called? Are any of the menagerie of non-classical logics, such as relevant logics, intuitionistic logic, paraconsistent logics or quantum logics, as deserving of the title 'logic' as classical logic? From another angle (or from another continent): Is classical logic really as deserving of the title 'logic' as relevant logic (or any of the other non-classical logics)? If so, why so? If not, why not?

Furthermore, there are debates *within* the camp of those who favour one logic, as to how this logic is to be understood. In the case of classical first-order logic, proponents of a Tarskian analysis of consequence say that logical validity amounts to preservation of truth-in-a-model. Others argue that this is not an adequate analysis of the notion of logical consequence and that we must look elsewhere for an adequate analysis of the notion [16, 17, 19, 54, 58]. How are we to understand this debate?

Logic has a chief subject matter: Logical Consequence. The chief aim of logic is to account for consequence — to say, accurately and systematically, what

[1] Except where grammar dictates otherwise, 'Logic' (uppercase 'L') names the discipline, 'logic' (lowercase 'l') names a system within Logic.

consequence amounts to, which is normally done by specifying which arguments (in a given language) are *valid*. All of this, at least today, is common ground.

Logic has not always been seen in this light. In older days, Logic was dominated by the Frege–Russell picture which treats logical truth as the lead character and consequence as a mere supporting role. The contemporary, post-Tarskian picture reverses the cast: *consequence* is the lead character; logical truth is a side-show. For example,[2] Etchemendy writes:

> Throughout much of this centry, the predominant conception of logic was one inherited from Frege and Russell, a conception according to which the primary subject of logic, like the primary subject of arithmetic or geometry, was a particular body of truths: logical truths in the former case, arithmetical or geometric in the latter ... This conception of logic now strikes us as rather odd, indeed as something of an anomaly in the history of logic. We no longer view logic as having a body of truths, the logical truths, as its principle concern; we do not, in this respect, think of it as parallel to other mathematical disciplines. If anything, we think of the consequence relation itself as the primary subject of logic, and view logical truth as simply the degenerate instance of this relation: logical truths are those that follow from *any* set of assumptions whatsoever, or alternatively, from no assumptions at all. [16, page 74]

But what is logical consequence? What is it for a conclusion, c, to logically follow from premises P? Putting things another way, what is it for an argument from premises P to conclusion c to be valid?

There is a venerable tradition to which virtually all contemporary philosophers, and many, many logicians, subscribe. According to this tradition the nature of logical consequence — the nature of "real validity" — is captured in the following principle:

(V) A conclusion c *follows from* premises P if and only if any case in which each premise in P is true is also a case in which c is true. Or equivalently, there is no case in which each premise in P is true, but in which c fails to be true.

Here is one example of the use of this principle to define validity. The quotation is taken from Richard Jeffrey's text *Formal Logic: its scope and its limits.*

> Formal logic is the science of deduction. It aims to provide systematic means for telling whether or not given conclusions follow from given premises, i.e., whether arguments are valid or invalid ...
>
> Validity is easily defined:
>
> > A *valid* argument is one whose conclusion is true in every case in which all its premises are true.
>
> Then the mark of validity is absence of counterexamples, cases in which all premises are true but the conclusion is false.

[2] For a more lengthy discussion of the centrality of *consequence* in logic, see Chapter 2 of Stephen Read's *Thinking About Logic* [42].

> Difficulties in applying this definition arise from difficulties in canvassing the cases mentioned in it ... [20, page 1]

So, this account of logical consequence, which might be described in shorthand as: *truth preservation in all cases*, is a plausible way to give an account of logical validity.

However, (V) does not give us a *complete* account of logical consequence. (V) alone does not take us very far down the road towards constructing a *logic*. For *that* we need an accurate and systematic account of which arguments are valid. (V) by itself does not give us an account of the *cases* involved.

Jeffrey's last line is significant: "Difficulties in applying this definition arise from difficulties in canvassing the cases mentioned it." In this paper we present a view that takes such "difficulties" very seriously. The view is *logical pluralism* — 'pluralism', for short. Pluralism, we believe, makes the most sense of contemporary Logic.

2 Pluralism in Outline

Pluralism, stripped down, is captured in the following list of conditions.

1. The pretheoretic (or intuitive) notion of consequence is given in (V).
2. A *logic* is given by a *specification of the cases* to appear in (V). Such a specification of cases can be seen as a way of spelling out *truth conditions* of the sentences in the language in question.
3. There are *at least two* different specifications of cases which may appear in (V).

Point (1) is self-explanatory. We have already seen a statement like (V) being used by Jeffrey to introduce the notion of logical consequence. The idea that (V) is an account of logical consequence is by no means idiosyncratic. We will not attempt an extensive search of the literature, though evidence for the centrality of an analysis like (V) is not hard to find. Here is one more case: W. H. Newton-Smith, in his popular introductory text, writes:

> [Some arguments] have true conclusions whenever they have true premises. We will say that they are *valid*. That means that they have the following property: In any case in which the premise (premises) is (are) true, the conclusion must be true. [35, page 2]

For us, logic is a matter of preservation of truth in all cases. We take this to be the heart of logical consequence.

However, this is not the end of the matter. To construct a *logic* you need to spell out what these *cases* might be. To give a systematic account of logical validity, you need to give an account of the cases in question, and you need to tell a story about what it is for a claim to be *true in a case*. Without an answer to these questions, you have not specified a logic.

This truism is given in point (2) of our account of logical pluralism. To use (V) to develop a *logic* you must specify the cases over which (V) quantifies, and you must tell some kind of story about which kinds of claims are true in what sorts of cases. For example, you might give an account in which cases

are *possible worlds.* (Furthermore, you might go on to tell a metaphysical story about what sorts of entities possible worlds might be [24, 25, 55, 62] — but then again, you might not.) On the other hand, you might spell out such cases as set-theoretic constructions such as *models* of some sort. However this is done, it is not the sole task. In addition, you must give an account of *truth in a case.* Here is an example of how you might begin to spell this out. Your account of cases and truth in cases might include this condition:

- $A \wedge B$ is true in x if and only if A is true in x and B is true in x.

where A and B are *claims*[3] and x is a *case.* Such an assertion tells us that a conjunction is true in a case if and only if both conjuncts are true in that case. This gives us an account of *truth* in cases which not only tells you how conjunction works, but it also gives you some data about validity. Once we have this connection, we have the validity of the argument from $A \wedge B$ to A. For any case x, if $A \wedge B$ is true in x then A is true in x, by the condition given above. This is but one example of how you might begin to systematically spell out the conditions under which claims are true in cases. To do this is to *do logic.*

None of this so far is particularly controversial.[4] The controversy in our position comes from point (3). According to the third and final claim, there are *different* ways to specify the "cases" appearing in (V). There is no *canonical* account of cases to which (V) appeals. There are different, equally good ways of spelling out (V); there are different, equally good *logics.* This is the heart of logical pluralism.

We will begin our elaboration and defense of (3) by examining different ways (V) can be — has been — filled out. We start with a very well-known way of filling out (V): Models for classical first-order logic.

3 Tarskian Models, and Classical Logic

There are many ways in which you might give an account of (V) which renders valid all of the theses of classical logic. One way is to treat the cases of (V) as *possible worlds.* Then your clauses for truth in a case, or truth in a world, will look like this.

- $A \wedge B$ is true in w iff A is true in w and B is true in w.
- $A \vee B$ is true in w iff A is true in w or B is true in w.
- $\sim A$ is true in w iff A is not true in w.

It is a little harder to give an account of the truth of quantified claims in possible worlds, but if we allow each object in each world to have a *name* in our language, then the clauses are trivial.

[3] Read our neutral 'claim' as picking out sentences, propositions, utterances, statements or anything else you think might feature in the premises and conclusions of arguments.

[4] Well, that is a small fib. One part *is* controversial. We have unashamedly priveleged the *model-theoretic* or *semantic* account of logical consequence over and above the *proof-theoretic* account. We think that a version of pluralism can be defended which does not privelege "truth in a case" to the same extent. However, since most of the current debates with which we are interacting lie firmly within this model-theoretic tradition, and since we are comfortable with that tradition, we are developing pluralism in this way.

- $\forall x A(x)$ is true in w if and only if for each object b in w, $A(b)$ is true in w.
- $\exists x A(x)$ is true in w if and only if for some object b in w, $A(b)$ is true in w.

Now, with no further analysis of what a world w might *be*, or how *many* there might be, a story of consequence can be told. We have already seen that this account validates the inference from $A \wedge B$ to A. It also validates the inference from A to $A \vee B$, from $A \wedge (B \vee C)$ to $(A \wedge B) \vee C$, from $\forall x(A \vee B)$ to $\forall x A \vee \exists x B$, and many more besides.

If the cases in our account encompass *all* possible worlds, then an argument is valid if and only if in any world in which the premises are true, so is the conclusion, or equivalently, if it is impossible for each premise to be true but for the conclusion to not be true. Call this the *necessary truth preservation* account of validity. This is one way to elaborate (V), but it is not the only one. In fact, it is not at all the *traditional* picture of logical consequence. The possible worlds account is not *formal* because it makes no essential use of the *forms* of the sentences analysed. To be sure, our elucidation has picked out conjunctions, disjunctions, negations and quantifiers, but there was no need at all to do this. We could just as well have given clauses for colour terms.

- *a is red* is true in w if and only if a is red in w.
- *a is coloured* is true in w if and only if a is coloured in w.

This explains why the necessary truth preservation account of validity renders the argument from *a is red* to *a is coloured* valid. It is valid because in any case (that is, in any possible world) in which something is red, it is also coloured. It is impossible that something be red and for it to fail to be coloured.

This is not the only way to account for logical consequence, and, as we have mentioned, it is not the mainstream tradition. According to logical orthodoxy, the argument from *a is red* to *a is coloured* is invalid, because it is not *formal*. It does not exploit any logical form: it has the form $Fa \vdash Ga$, and this form is invalid. We can give an account of this form of validity by varying the cases over which (V) quantifies. Now validity is a matter of form, and cases interpret *formal* languages. In our case, the languages of first-order logic, in which we have simple predicates, names, variables, quantifiers and propositional connectives. Then sentences in such a formal language are interpreted in a *model*. These are *Tarskian* models of first-order logic. A Tarskian model, M, is a structure that comprises the following:

1. A *nonempty set* D, the *domain*; and
2. A *function* I, the *interpretation*, satisfying the following conditions:
   (a) $I(E)$ is an element of D, if E is a name (in the given language);
   (b) $I(E)$ is a set of ordered n-tuples of D-elements, if E is an n-place predicate.[5]

[5]For the degenerate cases when $n = 0, 1$, we take a 0-tuple to be a truth value, and a 1-tuple to be an element of D.

Then we use a model to interpret sentence of the language. (We use assignments of values to variables, in order to interpret sentences with free variables. If α is an assignment of values to variables, $\alpha(x)$ is the value of the variable x. Furthermore, an x-variant of α is an assignment which agrees with α in the values of all variables except possibly x.)

- If α is an assignment of D-elements to variables, then $I_\alpha(x) = \alpha(x)$. If a is a name, $I_\alpha(a) = I(a)$
- $Ft_1 \cdots t_n$ is true in M, α if and only if $\langle I_\alpha(t_1), \ldots I_\alpha(t_n)\rangle \in I(F)$.
- $A \wedge B$ is true in M, α iff A is true in M, α and B is true in M, α.
- $A \vee B$ is true in M, α iff A is true in M, α or B is true in M, α.
- $\sim A$ is true in M, α iff A is not true in M, α.
- $\forall x A$ is true in M, α iff A is true in M, α' for each x-variant α' of α.
- $\exists x A$ is true in M, α iff A is true in M, α' for some x-variant α' of α.

We take models to be cases, and we have defined truth in a model, for sentences of a formal language, by the standard recursive clauses. This account then tells us about validity for arguments in the formal language, by way of (V). An argument is valid if and only if in every model in which the premises are true, so is the conclusion. For arguments of our natural language, validity is inherited by way of *formalisation*. We can define truth-in-a-model for sentences of English by the standard processes of regimentation of those sentences, and therefore we can define *validity* for natural language arguments. Call this account the *Tarskian* account of validity of arguments in natural language.

We now have our first dimension of plurality. Consider the question: Is the argument from *a is red* to *a is coloured* valid? We have seen that the answer is *yes* for validity as necessary truth preservation. The answer is *no* for the Tarskian account of validity. This argument has the form $Fa \vdash Ga$, and there are many models in which the premise is true and the conclusion false.

So, we have at least two different accounts of validity. One might now wonder: Is there any basis upon which to choose between these two accounts? Is there any reason you might prefer one to the other? The answer here is a resounding *yes*. Tarskian validity is *formal*; necessary truth preservation is not. Tarskian validity can (perhaps) be known *a priori*, but necessary truth perservation (probably) cannot. If Kripke is correct [22], the argument from *a is water* to *a is* H_2O is necessarily truth preserving, but this cannot be known *a priori*.

On the other hand, validity as necessary truth preservation does not rely on a choice of the family of logical consequence. Colour connections, temporal, spatial and other modalites, part–whole relations, and many other forms of necessary connections are equally encompassed by this account. The Tarskian account, on the other hand, makes a choice of logical constants, the priveleged parts of language which can contribute to logical form, and hence, logical validity. Not all Logic is simply a matter of form. (This, we think, is one major part of Etchemendy's criticism of the Tarskian account of logical validity [17]. It is not his *only* problem with this account.)

A pluralist on the question of *formality* will call both accounts *logic*. *Thoroughgoing* plurists, like ourselves, will be happy to call the result of both the Tarskian account, and the necessary truth preservation account, *logic*, for both are ways of spelling out the pretheoretic account (V) of logical consequence. The proper answer to the question 'is the argument from *a is red* to *a is coloured* ***really valid***?' is to say '*yes*, it is necessarily truth preserving, and *no*, it is not valid by first-order logical form.'

Our pluralist account of disagreement about logical form goes as follows: It's not *fruitful* to debate which of these things is *logic*. Both are a way of fleshing out (V), so both are *logic*. Given an argument which is necessarily truth preserving but not Tarski-style valid, it is surely more informative to say: yes, there is no possibility in which the premise is true and the conclusion false, but there is a Tarski-style model in which the premise is true and the conclusion false, and this shows the necessary truth preservation is not in virtue of the first-order logical form of the sentences involved. *That* is an informative analysis. A debate about which of these is *logic* has nothing to add to this.

However, this is not the only kind of problem people might have with the traditional Tarskian analysis of logical consequence and first-order logic. Consider the zero-premise arguments to conclusions such as $\exists x(x = x)$ or to $\exists x(Fx \vee \sim Fx)$. If cases comprise Tarskian-style models these arguments are valid. Famous debates have raged over this result. A long and rather formidable tradition claims that neither $\vdash \exists x(x = x)$ nor $\vdash \exists x(Fx \vee \sim Fx)$ is Really Valid; *logic*, in this tradition, allows for the empty case, but Tarskian-style cases are never empty — the domain of a Tarskian-style model cannot be the empty set.[6]

Similarly, there is another famous worry involving the so-called ω-rule. Suppose that F holds of every (natural) number n. Then each of $F(0)$, $F(1)$, $F(2)$, ... $F(n)$, ... is true. But, then, $\forall n F(n)$ — or $\forall x(N(x) \supset F(x))$ — is true. But, given the compactness of classical first-order consequence, $F(0), F(1), \ldots \vdash \forall n F(n)$ is not valid. But many have argued, strenuously, that it *is* valid. What should be said in this case (or the case above)?

For now, set the details of these debates aside. The foregoing worries fit an important pattern, a pattern to which we will return in later sections. Roughly, the pattern is this. Let $C(L)$ be an account of consequence, or some precisification of 'validity', where L is the corresponding logic. Then $C(L)$ is said to *undergenerate*, with respect to some argument, if that argument is Really Valid but not $C(L)$-valid. The problem, in this case, is that $C(L)$ gets things wrong by failing to call the argument 'valid' when "in fact" it is valid — *Really Valid*. $C(L)$ is said to *overgenerate*, with respect to some argument, if the argument is *not* Really Valid but *is* $C(L)$-valid. In this case, $C(L)$ gets things wrong by calling the argument 'valid' when "in fact" it is not.

The undergeneration–overgeneration pattern is ubiquitous in philosophy of logic; indeed, it may well be *the* central pattern of dispute in the field.[7] The

[6] Another famous objection (which we will not pursue at this point) is voiced by Kreisel, Boolos, and McGee [9, 21, 28], to the effect that the models given in the traditional Tarskian account of validity are too limited. (Why not allow for domains too "big" to be sets? *Logic*, alone, seems not to impose this restriction, but traditional Tarskian-style cases do.)

[7] Maybe. We take no firm stand on this point here, though we are interested in this part of history. (The terminology involving 'undergeneration' and 'overgeneration' here is found in Etchemendy [17], Sánchez-Miguel [53], O'Hair [36], Read [41, 42], and others.)

important point here is that pluralism can make sense of the debate, though, in general, it refrains from blessing *one* side of the debate with the honorific title 'logic'.

A pluralist response to these issues goes as follows: Many appeals to "Real Validity" are appeals to *real* validity; they are not, however, appeals to the *only* real validity. Real validity comes from a specification of cases which appear in (V). According to pluralism, there are at least two such specifications of cases. Thus far, we have looked at two different approaches within *classical* logic — the worlds approach, and the Tarskian models approach. But these are just the beginning.

4 Situations, and Relevant Consequence

Each of the accounts of interpretations or truth conditions seen so far have been *classical* with respect to negation. For any cases x seen so far, be they worlds, Tarskian models, class-size models, or even models with empty domains,

- $\sim A$ is true in x if and only if A is not true in x.

Call this the *classical negation clause*. More properly, it should be called the classical negation clause *for x's*, since its content depends essentially in what kind of case x might be.

There are many good reasons for using a classical negation clause in constructing an account of truth in cases. The most obvious reason is the way that we use negation, and the conditions under which negations are, in fact, true: $\sim A$ is true just when A is not true. This is simply what 'not' *means*.[8]

However, to infer from this truism that the classical negation clause is the only acceptable reading of negation in any kind of *case* worth using in elaborating (V) would be far too swift. To do so would be to assume that the only acceptable use of cases is to model (or to *be*) consistent, complete *worlds*. Many have questioned this assumption. There are other ways to give an account of cases, or conditions under which claims might be true or false, which do not restrict themselves to consistent complete cases. One such line of account is the from the *situation theory* of Barwise and Perry [1, 2, 3, 11].

For situation theorists, the world is made up of situations. They are simply *parts* of the world. Claims are true of not only the world as a whole, but some claims at least are true of situations, these particular parts of the world. We do not need to spend time on the theory of situations and their individuation here: we need simply illustrate it with some examples. In the situation involving Greg's household as he writes this, it is *true* that Christine is reading a paper. It is also *true* that the stereo is playing. It is *false* that the television is on. It follows from this, and the fact that the television is in fact an inhabitant of the situation, that it is *true*, in this situation, that the television is off.

Situations "make" some claims true, and they "make" others false. However, some situations, by virtue of being *restricted* parts of the world, leave some claims undetermined.[9] It is not true in this situation that JC is reading. It is

[8] You might well say, instead, that this is what *true* means.

[9] We use shudder quotes around 'make' here not that we wish to avoid the use of "truth-making" terminology. To the contrary, we value the recent revival of this terminology and the analysis of the connections between claims and parts of the world which make them true [18, 33, 47]. We know, however, that this terminology is not used by situation theorists, and that it would be a mistake to impute it to them.

also not false in this situation that JC is reading — that is, it is not true in this situation that JC is *not* reading. JC does not feature in this situation at all.

It follows that the classical account of negation fails *for situations*. This treatment of negation is out of place in this context. It seems plausible, however to hold fast to the classical analyses of conjunction and disjunction.

- $A \wedge B$ is true in s if and only if A is true in s and B is true in s.
- $A \vee B$ is true in s if and only if A is true in s or B is true in s.[10]

We must emphasise that this point that the non-traditional treatment of negation does not mean that we are modelling a *non-classical* negation. Quite to the contrary. Our treatment of negation on this topic is not the traditional one simply because we are entering a new field — the logic of situations. It has not been traditional to formally model claims of the form 'A is true in situation x'. Once you do so, and once you acknowledge that situations are restricted parts of the world, it becomes clear that you ought reject the classical treatment of negation when applied to situations. This is completely consistent with the classical treatment of the truth or falsity of negation *simpliciter*. We may maintain that $\sim A$ is true if and only if A is not true. That is not in question. The situation theoretic analysis of this equivalence will proceed further: $\sim A$ is true if and only if $\sim A$ is true in some (actual) situation or other. A is not true if and only if A is not true in any (actual) situation whatsoever. The traditional, classical equivalence is maintained if we agree, then, that if $\sim A$ is true in some (actual) situation, then A is not true in any (actual) situation. And this is simple to maintain, given three, plausible, theses.

- There is a situation w, of which every actual situation is a part.
- If A is true in s and s is a part of s' then A is true in s'.
- If s is an actual situation, and if $\sim A$ is true in s then A is not true in s.

These minimal theses connecting negation and situations ensure the truth of the classical account of negation (simpliciter). Negation, here, is classical.

One of us [GR] has argued elsewhere that a similar case can be made for understanding a version of Łukasiewicz's three-valued logic of future contingents [26, 43]. Claims can be made true by history up to this moment, or made false by history up to this moment, or left unmade by history up to this moment. This is consistent with the *entire* history being classical with respect to negation. Either there will be a sea battle tomorrow, or there will not. *Tomorrow* will see to that. *Today*, however, sees to neither disjunct, and hence neither to their disjunction.[11]

We have not, however, given a systematic treatment of the truth or falsity of negations in situations. This can be done in any of a number of ways. You can,

[10] The conjunction clause is never disputed. The disjunction clause sometimes is, though we think that its rejection is, by and large, ill-motivated. If in this situation the milk is on the table or in the fridge, then either in this situation the milk is on the table or in this situation it is in the fridge.

[11] This differs from, but is consistent with, a supervaluational account of truth-at-a-moment, which identifies this with truth in all histories passing through this moment [4, 60]. Some things are true now, not in virtue of what exists now, but in virtue of what will happen, no matter how things turn out — such as tomorrow's sea battle or its absence.

for example, take satisfaction and dissatisfaction of relations in situations as primitive, and then inductively build up truth and falsity conditions of complex sentences.[12] This approach is traditional in situation theory, and it is also used in some varieties of semantics for non-classical logics [2, 5, 13, 34, 46]. Here, however, we will favour a different approach — the *compatibility* approach, which stems from Dunn's analysis of the semantics of negation [14, 15, 48]. On this proposal negation is taken to act in situations rather like necessity or possibility does in possible worlds. We admit into our semantics non-actual situations (or *models* of non-actual situations) which are connected by a binary relation of *compatibility*, which we write 'C'. Given this apparatus, we need not treat truth and falsity in parallel — negation is definable.

- $\sim A$ is true in s if and only if A is not true in s' for any s' where sCs'.

The negation $\sim A$ is true in s just when any situations in which A is true are incompatible with s. This clause follows fairly immediately from the meanings of negation and compatibility. If $\sim A$ is true in s and A is true in s', then s is not compatible with s'. Conversely, if A is not true in any s' compatible with s, then it appears that s has *ruled A out*. That is, $\sim A$ is true in s. This reading does not rely on a "funny" negation; it is completely compatible with a classical view of negation.[13]

Given such a semantics of situations, a natural reading of (V) emerges: a *situated* reading.

> The argument from P to c is *relevantly* valid if in any situation model, in any situation in which all premises in P is true, so is c.

This has an obvious sense: any situation in which all of the premises are true is one in which the conclusion is also true. To speak loosely but suggestively: To make the premises true you make the conclusion true too. The *relevance* of this reading of consequence is immediate. The inference from A to $B \vee \sim B$ fails, since a situation in which A is true need not be one in which $B \vee \sim B$ is true.

If we take the relevant *tautologies* to be those claims true in every situation, then $B \vee \sim B$ is not among them. This does not mean that we have adopted a strange non-classical account of negation. We agree with the classical theorists that $B \vee \sim B$ is true, that it is true in *every* world. Our negation is classical. The invalidity of the argument from A to $B \vee \sim B$ is a *relevant* invalidity. (And this is as it ought to be. $B \vee \sim B$ is true, after all, but it need not follow *from* the truth of an arbitrary A, at least when we take the 'from' seriously.) The argument, of course, is still *classically* valid, in the sense that in any world in which A is true, $B \vee \sim B$ is true.

The move to situations as incomplete parts of the world is a natural one. It has a natural generalisation, which is also available to us, to consider not only *incomplete* situations, but also *inconsistent* situations, or *ways things could not*

[12] For example, you will say that not only is a conjunction true in s when both conjuncts are true, but dually, a conjunction is false when one conjunct is false.

[13] The three minimal conditions cited earlier for a classical treatment of negation have their "compatibility" readings. (1) Any actual s is a part of a world w (this is as before); (2) w is a world if and only if wCw, and if wCs then s is part of w (in other words, worlds are maximal, self-compatible situations); (3) if sCt, s' is a part of s and t' is a part of t, then $s'Ct'$ too (compatibility of wholes leads to compatibility of parts).

be [30, 48, 49, 52, 63]. These are situations which fail to be self-compatible. If s is not compatible with itself, then it is possible that both A and $\sim A$ be true at s. This, again, is not terribly non-classical. According to our account of worlds as consistent, complete situations, such *impossibilia* cannot be a part of any *world*. Worlds are consistent, and hence, have no inconsistent parts. This does not mean, of course, that there are no *ways that things could not be*; it means, simply, that the worlds are not (and could not be) among them.

Given the admission of inconsistent situations, argument from $A \wedge \sim A$ to B fails the relevant test, for a situation in which $A \wedge \sim A$ is true need not be one in which B is true. A situation might well be *inconsistent* about A without involving *everything*. This same situation gives us a counterexample to the relevant validity of *disjunctive syllogism*: the argument from $A \vee B$ and $\sim A$ to B. A situation inconsistent about A but not judging B as true suffices. $A \vee B$ is true in this situation, as is $\sim A$, but B fails.

This last case has been the cause of much debate in the literature on relevant logics and relevant inference. Much ink has been spilled on the failure of disjunctive syllogism and whether it is a virtue or a vice [29, 40, 51]. We do not plan to add to the spilling of ink in any depth here. We will simply note that traditional criticisms of the relevant rejection of disjunctive syllogism are beside the point, when seen in the light of pluralism. We will end this section on relevant consequence by explaining why this is so.

One cause of concern with the rejection of disjunctive syllogism is that disjunctive syllogism is *obviously* valid, and we reason with it all the time — we could not do without it in everyday reasoning [6]. The pluralist can agree: *Of course* there is a sense in which disjunctive syllogism is valid — and even *obviously* so. In any possible world (or Tarskian model) in which the premises are true, so is the conclusion; it is impossible for the premises to be true but the conclusion be false. In *that* sense — the sense afforded by cases as world or as Tarskian models — disjunctive syllogism is obviously valid. The virtue of a pluralist account is that we can enjoy the fruits of relevant consequence as a guide to inference without feeling guilty whenever we make an inference which is not relevantly valid. With classical consequence you know you will not make a step from truth to falsehood. With relevant consequence, the strictures are tighter; you know you will not make a step from one that is true in a situation to something not true in it (but which might be true outside it). This is a tighter canon to guide reasoning.

So, the case of incomplete and inconsistent situations motivates a genuinely different elucidation of logical consequence — one which differs with the classical account on the validity of inferences down to the propositional level. This account of consequence is still recognisably *logic*; it is another way to flesh out our condition (V). It is not a *rival* in any sense to the classical, traditional explications of that condition. Instead, it can coexist alongside classical validity as *another* important variety of logical consequence.[14]

[14] We have restricted our attention here to the conjunction, disjunction and negation fragment of relevant logics. More can be done to bring the notion of relevant entailment *into* the language at hand. For another approach to relevant logics which motivates *two* varieties of consequence, but from a very different perspective, we refer the reader to Mark Lance's "Two Concepts of Entailment" [23].

So now we have even more pluralism, a pluralism in which (V)'s cases may be worlds, Tarskian models, incomplete situations, and even incomplete and inconsistent situations. Before moving on, we should perhaps clarify a potential confusion.

One might read what we've been saying as one big *ontological* (or perhaps *epistemological*) thesis: "*If* reality *really does* comprise situations, then of course we'd better use relevant consequence! Or *if* reality comprises inconsistent situations, then we'd better use paraconsistent consequence! Or *et cetera*." In reply to this potential reading, let us say briefly but clearly: *No!* — at least not here. Pluralism is neither an ontological nor an epistemological thesis. We are talking about Logic, and about (V); we are saying that situations can serve as (V)'s cases just as much as worlds, Tarskian models, or whatever else you might have in mind. Whether such situations, worlds, Tarskian models *et cetera* Really Exist, or whether (or in what sense) we Really Know about them, are questions to which we remain neutral, at least in this paper. To be sure, some of the ways of specifying (V)'s cases may well be motivated by various ontological sorts of intuitions — intuitions with respect to what Reality is like, how language and the world get together, and so on. (One obvious example, here, may be possible worlds.) Though we find such issues to be interesting, they remain issues for another occasion. In the end, we take it to be a separate question as to how (V)'s quantifiers should be read, whether they carry "existential import", how we come to know about (V)'s various cases, or *et cetera*. For present purposes, we remain neutral with respect to such questions.[15]

We will return to this sort of worry in a later section. (See, specifically, Section 6, "Comments, Criticisms, and Questions".) For now, turn to another way in which (V)'s cases can be, and have been, filled in.

5 Constructions, and Intuitionistic Logic

Mathematicians do not, generally speaking, concern themselves with a *situated* account of logical consequence while reasoning about mathematical objects or structures. However, they too can make some distinctions which are blurred by classical accounts of validity. We have in mind the mathematics pursued by mathematical *constructivists*.

The constructivism of the mathematicians Errett Bishop [7, 8], Douglas Bridges [10], Fred Richman [32, 50] and others can best be described as mathematics *pursued in the context of intuitionistic logic*.[16] In constructive mathematics the goal is to gain understanding of mathematical structures, and to prove theorems about them (just as in classical mathematics). However, the goal is to prove mathematical theorems with constructive, or computational content. If a statement asserting the existence of some mathematical object is proved in a constructive manner (using the rules of intuitionistic logic) then this proof will contain the means of specifying the object or structure in question. Wittgenstein illustrates the advantages of constructive proof over its classical

[15] Of course, there's room for different sorts of pluralism so understood. For example, one sort of pluralism might be an *ontological* pluralism, one that recognises all manor of existents corresponding to the plurality of (V)'s cases. One might also be an *epistemological* pluralist, taking there to be genuinely different ways of knowing about (V)'s plurality of cases, or *et cetera*. We leave these issues for another occasion.

[16] Tait provides more explicitly *philosophical* account which draws very similar distinctions to the work of constructive mathematicians [56, 57].

cousin by drawing out its implications for our *understanding*.

> A proof convinces you that there is a root of an equation (without giving you any idea *where*) — how do you know that you understand the proposition that there is a root? [61, page 146]

This feature of constructive mathematics is guaranteed by the structure of constructive proofs. We emphasise the fact that this is a new notion of proof by using the word 'construction' for this notion. Constructions obey the following laws:

- A construction of $A \wedge B$ is a construction of A together with a construction of B.
- A construction of $A \vee B$ is a construction of A or a construction of B.
- A construction of $A \supset B$ is a technique for converting constructions of A into constructions of B.
- There is no construction of $\bot$.[17]
- A construction of $\forall x A$ is a rule giving, for any object n, a construction of $A(n)$.
- A construction of $\exists x A$ is an object n together with a construction of $A(n)$.

This elucidation is *not*, at its heart, formal. It is an informal account because it leaves its central notion (that of a construction) undefined.[18] This notion can be formalised in many different ways — we will see one account later in this section. For all its informality, however, it gives us an understanding of the behaviour of constructive proof. For example, the inference from $\forall x(A \vee B)$ to $\exists x A \vee \forall x B$ is clasically valid, but not constructively acceptable. For example, it is easy to demonstrate that every string of ten digits in the decimal expansion of π is either a string of ten zeros, or it is not. A demonstration of this fact, however, does not give us a construction of the claim that either there is a string of ten zeros in π or every string of ten digits in π is not a string of zeros. Any construction of *this* statement must either prove that there is no string of ten zeros in π or to show where one such string is. The constructive content of $\exists x A \vee \forall x B$ is greater than that of $\forall x(A \vee B)$.

Theorems of constructive mathematics are simply theorems of mathematics proved construcively. According to this approach, the theorems of constructive mathematics are also theorems of classical mathematics. The difference between constructive and classical mathematics is not one of subject matter, but one of the required standards of proof. Classical mathematicians may appeal to the law of the excluded middle, and proof by contradiction. Constructive mathematicians do not, for these moves destroy constructivity.

It might be thought that intuitionistic logic, defined in terms of constructions and proofs is less amenable to our "truth conditional" account of logical

[17] We define $\sim A$ as $A \supset \bot$, so a construction of $\sim A$ is a technique for converting constructions of A into absurdity. It shows that there are no constructions of A.

[18] See the similarity to the account of worlds at the start of Section 3. We gave an account of what it is for a conjunction to be true in a world. We gave no account of what it is for an arbitrary sentence to be true in a world. Similarly here, we give no account of what it is for an arbitrary statement to be given by some construction.

consequence given in (V). This is not the case. A truth conditional semantics may be given for the intuitionistic logic of constructive mathematics, which both does justice to the practice of constructive mathematics and opens the way for a pluralist reading of that practice. The truth conditional semantics is simply Kripke's semantics for intuitionistic logic. Truth is relativised to *points* (which model constructions) which are partially ordered by *strength* (written '$\sqsupseteq$').

- $A \wedge B$ is true in c if and only if A and B are true in c.
- $A \vee B$ is true in c if and only if A is true in c and B is true in c.
- $A \supset B$ is true in c if and only if for any $d \sqsupseteq c$, if A is true in d then so is B.
- $\sim A$ is true in c if and only if A is not true in d for any $d \sqsupseteq c$.

The points in a Kripke structure for intuitionistic logic do a good job of modelling *constructions* ordered by a notion of relative strength. The clauses for conjunction and disjunction are straightforward transcriptions of our pre-formalised notion of constructions. The rules for implication and negation differ somewhat, but can be motivated to follow from the pre-theoretic notion. A construction proves $A \supset B$ if and only if when *combined with* any construction for A you have a construction for B. The assumption guiding Kripke models is that a construction for $A \supset B$ combined with one for A will be a *stronger* construction.[19] So, $A \supset B$ is true at c if and only if any stronger construction d for A is also a construction for B.

Constructions are *incomplete* and hence should not be expected to construct, for every claim A, either it or its negation $\sim A$. Constructions have computational content, so a construction of $A \vee B$ should be a construction of A or a construction of B. This jointly ensures that $A \vee \sim A$ ought fail. This can not necessarily be constructed.

Of course, for a pluralist, it does not follow that $A \vee \sim A$ is not true, or even, not *necessarily* true. It is consistent to maintain that all of the truths of classical logic hold, and that all of the arguments of classical logic are *valid* with the use of constructive mathematical reasoning, and the rejection of certain classical inferences. The crucial fact which makes this position consistent is the shift in context. Classical inferences are valid, *classically*. They are not *constructively* valid. If we use a classical inference step, say the inference from $\forall x(A \vee B)$ to $\exists x A \vee \forall x B$, then we have not (we think) moved from truth to falsity, and we cannot move from truth to falsity. It is impossible for $\forall x(A \vee B)$ to be true and for $\exists x A \vee \forall x B$ to be false. However, such an inference *can* take one from a truth which can be constructed to one which cannot, as we have seen. So, the inference, despite being classically valid, can be rejected on the grounds of non-constructivity.

This pluralist account of constructive inference is not a view that will be shared by constructivists who wholeheartedly reject the use of classical inference. However, it is a classical view which makes sense of what constructivists are doing. According to this view, constructive mathematics retains its meaning. When a constructivist says 'not', she means 'not'. She does not mean something

[19] This is the assumption challenged by relevant accounts of implication. In constructive mathematics, where relevance is not at issue, this account is appropriate.

else, foreign to the classical mathematician.[20] She differs from the classical reasoner only in her use of tighter canons of inference. It is hard to see how any other view can do justice to the practice of constructive mathematics. It seems that classical dogmatists must either reinterpret constructivist claims as being about something else (when she says $\sim P$ she means not that P is not true, but instead that P can be *proved* not to be true) or that intuitionistic logic merely a formalist game in which the rules are *syntactically* restricted to allow a more limited repertoire of proof.

6 Comments, Criticisms, and Questions

We have given an account of logical pluralism, and we have shown how it contributes to our understanding of different traditions in contemporary Logic. In this section we will examine potential criticisms of the view.

Anything Goes?

Objection: "You are *pluralistic* about logical consequence. You say that there is no objective matter about whether a conclusion *follows from* premises. Does it not follow that *anything goes*? On your view, there is no disagreement about logical consequence. But that makes a mockery of the current state of play in logic. There is disagreement all over the place, when it comes to giving a correct account of logical consequence. For example, Stephen Read writes:

> Rival logical theories, such as intuitionistic logic, paraconsistent logics, relevant logics, connexive logics, and so on, are based on different philosophical analyses of this basic notion [logical consequence].
>
> [42, page 36]

"According to your view, these logics are not *rivals*, they live in one large happy family. Similarly, Graham Priest writes:

> Whether or not any of the nonstandard logics discussed here [viz., intuitionist, many-valued and quantum, relevant and paraconsistent, conditional and free] are correct, their presence serves to remind that logic is not a set of received truths but a discipline where competing theories concerning validity vie with each other. [39]

"On your account, these theories do not *compete*. The people who propose such logics certainly *intend* them to compete with other logics. You have misunderstood contemporary Logic."

Reply: This is a challenging criticism, but we think it is ultimately misguided. Pluralism is *not* a recipe for wholesale agreement. *Of course* there can be disagreements about logical consequences. Our pluralism holds that *some* formal logics can fruitfully be seen as different elicidations of (V), the pretheoretic notion of logical consequence, and that (V) does not determine *one* logic, but rather, a number of them. It does not follow that there are no disagreements

[20] We keep using the example of the mathematician, merely because constructive reasoning is most developed in this tradition. It need not be restricted to mathematics. Mathematical technique is *applied* when talking about the environment. We can reason constructively not only about the real line, but also about spatial and temporal distances, physical quantities, and many more things besides.

about notions of logical consequence. It does follow, however, that in any such disagreement the ground has to be fixed to ensure that the disputants are not talking past each other.

Let us give some examples of what we mean, by illustrating two positions with which we disagree, while maintaining our pluralist credentials. We disagree with the dialetheism of Graham Priest, according to which contradictions may be true. According to Priest, there are arguments of the form $A, \sim A \vee B \vdash B$ which are not only *invalid* but which have true premises and an untrue conclusion [37, 38]. We disagree.[21] So, this means that when we cash out validity as necessary truth preservation, or anything similar, we hold that all of the inferences of classical first order logic are *valid*. Priest does not.

In a similar vein, we also disagree with Dummett, and other intuitionists, who hold that there are arguments from A to $B \vee \sim B$ with true premises and untrue conclusion [12]. We disagree; we take every instance of $B \vee \sim B$ to be true, and to be true of necessity. These folk and we disagree.[22]

In these cases disagreement is possible. It is possible once we have set the terms of the debate. In both cases, with paraconsistent and intuitionist logic, we find a place for these non-classical logics — for both are elucidations of the pretheoretic notion (V) of logical consequence.

There might be other logics for which we can find *no* place for in our catalogue of True Logics, as much as we admire their technical subtlety. There are too many modal logics, for example, to hold each of them as the logic of broad metaphysical necessity. So, given a particular interpretation of each of the symbols in our formalism (including *consequence*, which is rarely spelled out) we can admit that there is a great deal of scope for rivalry. For the propositional modal logic of necessary truth preservation, we think that a logic somewhere between S4 and S5 is a candidate for getting things *right*. Anything else *gets it wrong* when it comes to metaphysical necessity, we think. So, there is scope for rivalry and disagreement, when the meaning of the basic lexicon is settled. The moral of our pluralism goes as follows: Once you are specific about what your logic is meant to be *doing*, there is scope for genuine disagreement. But before that is done, who knows what is going on?

This raises a general question: What is it to *disagree* with an account of consequence? What kinds of disagreement are possible? We think that there are at least four different ways in which disagreement and difference between formal logics can be understood. Here is a rough spectrum of what one might think about a logical system L.[23]

- *Abstract Geometries*: This is the analogy with geometry. L is a *logic* because it is formally similar to other logics. It models a consequence relation. It is to logical systems what a finite projective geometry is to Euclidean geometries. Euclidean geometries and their close neighbours are used to model physical space. A finite projective plane is not going to be

[21] Well, *here* we do. We have great respect for the dialethic tradition in logic (according to which there are true contradictions), and at times one of us has pursued it [44, 45].

[22] There is nothing essential to pluralism to take the classical side of the debate in these disagreements. We could just as well be pluralist dialetheists or intuitionistis. For example, we might hold that some propositions of the form $A \wedge \sim A$ were true, and still accept both a "classical" paraconsistent logic and a constructivist one, to model constructive reasoning.

[23] Thanks to Daniel Nolan for discussion on this point.

used to model physical space, but it might well be used to model something analagous to physical space. Simiarly, system L might be used to study something analagous to consequence relations between propositions. And so, it is called a logic for reasons of structural similarity.

- *Applied Geometries*: Take two geometries, say, a three-dimensional Euclidean space, and a particular non-Euclidean three-dimensional space. These two spaces might be *competing* models for the physical space in our region. Here the geometries are *applied*, for there is a notion of what it is to which the theoretical entities must correspond. Once rules of application of the model are settled, there is scope for a genuine *disagreement* between the two theories.

  Similarly, once applied, there is a scope for genuine disagreement between logical systems. However, we, as pluralists, do not think that this disagreement comes about *simply* by applying the logic to model the validity of real argument. We think that different formal systems can be equally appropriately used to model the validity of arguments. The analogy with applied geometry becomes appropriate only once the pretheoretic account (V) is fleshed out. Once you have a specific account of what kind of cases are in use (be they, worlds, constructions, situations, whatever) then there is scope for disagreement.

- *Different Subject Matter*: We do not know how to label this position. X thinks what Y is doing is attempting to get at the same kind of thing as what X is trying to get at, but Y is going about it in completely the wrong way, and is actually either doing gibberish or talking about something else. The intuitionist view of the classicalist, or vice versa, *can* be seen like this, but need not be. A debate between the two which hinges upon whether the *proper* analysis of meanings ought proceed by way of truth conditions or in terms of provability or evidence conditions can be seen in this way.

- *Pluralism*: Finally, you can hold that two different logics L and L' are *both* accurate and systematic accounts of the one notion of logical consequence. We hold that this position is the appropriate one in each of the cases we have discussed.

We think that *all* points on this spectrum are inhabited.

Yes, Yes, *But*...

Objection: "Yes, yes, we can all agree that there are different ways of spelling out (V), and so different senses of the expressions 'valid', 'follows from', and 'consequence'. *But*, of course, what we really want to know is which of the various senses of 'valid' — *classical validity*, *relevant validity*, or *etc.* — is The True Sense?"

Reply: This is a bad question. The point that pluralism makes is this: Do *not* admit that 'valid' is ambiguous and, then, ask the question "But *which* of these senses is The True Sense?" In general, this question just does not make sense. The only arbiter of "Real Validity" is (V) itself; on this, however, (V) is silent.

One True Logic After All?

Objection: "Another potential problem with pluralism comes from the other direction. You have shown that there is a number of different ways that 'case' can be interpreted in (V). But (V) has a *universal* quantifier in the front. It says that an argment is valid if and only if in *all* cases in which the premises are true, so is the conclusion. Is not *real* validity then preservation of truth across *all* cases? Will this not mean that the *true logic* is the intersection of all logical systems given by (V)? You have *one true logic* after all."

Reply: We think that this conclusion should be resisted for a number of reasons. First of all, we think that classical first order logic *is* logic after all. We think that if the premises of a classically valid argument are true, so is the conclusion. Those arguments are *valid*. True, they are not all constructively valid, or relevantly valid, but this does not stop them being valid, in an important and useful sense. The class of all Tarskian models is an important and natural class of cases, and it is appropriate to restrict our quantifiers in (V) to those cases.

Secondly, we see no place to *stop* the process of generalisation and broadening of accounts of cases. For all we know the only inference left in the intersection of all logics might be the *identity* inference $A \vdash A$. How bizarre to say that *identity* is the only *really* valid argument! It seems much more appropriate respect contemporary use of the term to call each of these systems *logics*.

Finally, each formal system is used to regulate inference, each falls under our original pretheoretical banner for logical inference. So, each of them are *logics*.

Nontransitive Consequence?

Objection: "You parade your pluralism not only as the best approach to logical consequence, but also as a framework in which to understand contemporary Logic. Trouble is, your pluralism does not clearly afford understanding of the whole scene. For example, what, if anything, does pluralism say about "logics" in which consequence is not transitive? Presumably, (V) yields only transitive accounts of consequence; if so, then the "logics" in question are not really *real logics* after all. But if these aren't logics, what are they? And if they *are* logics, how can they be — given that, again, (V) seems to yield only transitive accounts. (Regardless of how (V)'s cases are spelled out, (V) still yields only *preservation*-type accounts of consequence; and in this context, the requisite sort of preservation seems to be essentially transitive.)"

Reply: This is a very serious objection which, unfortunately, cannot be dismissed due to a misunderstanding of our presentation of pluralism. As the objection says, we maintain that Logic is in the business of generating consequence relations, and that such relations arise from specifications of (V)'s cases. But, as the objection notes, (V) yields accounts in which consequence amounts to *preservation* — necessary truth preservation for classical logic, necessary situated truth preservation for relevant logics, and necessary "construction–proof" preservation for constructivist mathematics. In a significant sense, varying the logic is a matter of varying *how much* is to be preserved.

But the point of the objection concerns an apparent corollary: Pluralism, given essentially in terms of (V), seems to have it that consequence must be reflexive and transitive — no amount of varying cases can vary this. But, then, as the objection states, our pluralism rules out some activity in contemporary Logic. For example, the Martin and Meyer system S-for-Syllogism, which rejects $A \vdash A$ on grounds of circularity, is ruled out given the lack of reflexivity [27, 31].

Moreover, Tennant style "relevant logics" which reject transitivity likewise fail to fall under the banner of logical consequence given in (V) [59].

What can we say? Our opinion on this is not settled. However, our initial response is to maintain the line. These systems are appropriately called 'logics' purely by analogy with logics *properly* so called. Non-transitive or non-reflexive systems of "entailment" are perhaps modelling interesting phenomena, but they are not modelling *logical consequence*. One must draw the line somewhere.

Why Truth Rather Than Assertibility?

Objection: "Why fix on preservation of *truth* for logic? There are folks who think that preservation of warranted assertibility (or the like) is what's crucial. How, if at all, does pluralism make sense of *those* so-called logics?"[24]

Reply: Fans of assertibility say that logic, at least in the first instance, is a matter of preserving warranted assertibility rather than truth. But pluralism can respect this. Pluralism says that (V) captures the essentials of logical consequence. But the essentials *need* not be expressed with 'truth'. For the assertibility folks, the core of logical consequence can be rephrased slightly thus:

(V′) A conclusion C follows from premises $P_1, \ldots, P_n$ iff any case in which P_1 and P_2 ... and P_n, C.

The idea is that an argument $A \vdash B$ is valid iff any case in which A is a case in which B — any case in which A, B. Here, of course, we are *using* 'A' and 'B' instead of mentioning them. The important point is that, according to pluralism, there's room in the spectrum of (V)'s cases for warranted assertibility. This still affords the "truth-functional" account we have privileged in this paper; however truth, itself, need not be involved.

Logical Truth and Pluralism

Objection: "You cannot make sense of *logical truth*. Logical truths are those *truths* that are consequences of no premises whatsoever. But these are indeed *truths*; they're *not* "truths in X" for the various cases, X, you've discussed."

Reply: We can allow that 'logical truth' is ambiguous; indeed, we can allow that 'logical truth' has as many senses as there are ways of cashing out (V). Compatible with all this, however, is a view according to which all logical truths, in whatever sense, are in fact *true*. Nothing in pluralism conflicts with this, or requires it.

Of course, a sensible *constraint* on (V) might be that *one* of the cases *be* the actual world, or that one of the cases *model* the actual world. Once this constraint is imposed the resulting logical truths are true. But again, nothing in (V) seems to *require* this constraint.

Logical Consequence is more than (V) alone

Objection: "I agree that (V) is an important component of our notion of consequence. I think it is even arguable that it is the core of our intuitive or pretheoretical notion of consequence. And I think that the argument towards

[24] Thanks to James Chase for pressing us on this matter.

the end of the paper that taking (V) to be a common core of Logic helps explain logical practice gives us good reason to suppose that (V) is the core of consequence, or maybe even that our notion of consequence is in practice, or should in practice be, no more than (V). But (and here's the rub) I don't see any reason to say that it is our *entire* intuitive notion.

"But it is not plausible that (V) is the only platitude about logical consequence (though it seems to me that if (V) gave the *entire* intuitive notion of consequence, something like this might have to be true). Here are some platitudes about consequence:

1. Logical consequence never holds between a set of premises which are true simpliciter and a conclusion which is false simpliciter.

2. Logical consequence has something to do with good argument.

3. Logical consequence is somewhat like the thing that the people we call logicians (or lots of logicians, or most logicians) are formalising, or attempting to formalise.

4. Logical consequence (or deductive logical consequence, at least) is usefully formalisable, at least to a great extent.

5. The relations of logical consequence involving "and", "or" and "not" have close connections with the meanings of those connectives.

6. A lot of useful information about logical consequence is available in principle *a priori* — or at least, without needing too much particular investigation of specific *a posteriori* subject matters.

7. The logical consequences of $A \wedge B$ are not in general all the same as the logical conseuqences of $A \vee B$, the logical consequences of $\sim A$ are in general not all the same as the logical consequences of A.

"These are chosen at random, and hopefully they will seem intuitive to most people. Now, nothing you say means you can't take them to be intutive either, but if that's right, then it follows that (V) is not our *entire* intuitive notion of consequence."[25]

Reply: This is an important objection, and it is one to which we are very sympathetic. We allow that there may be *other* constraints to what counts as a logic, but we also maintain that this does not pin down the field of play to just one. That is the minimal requirement for pluralism.

However, it is interesting to note how many of the platitudes given above actually stem from (V).

1. If the actual world is (or is modelled by) a case, then indeed there is no valid argument from true premises to a false conclusion.

2. If for an argument to be *good* you do not step from truth to falsity (in any *case*) then of course logic has something to do with good argument.

[25] Thanks to Daniel Nolan for forcefully voicing this objection. Our presentation of this objection is a small rewriting of an email exchange with him.

3. Logicians themselves say that logic is about truth preservation in cases, so this is what logicians study.

4. We have seen in all of our elaborations of different accounts of cases *how* logic is formalisable, at least in the simple language of conjunction, disjunction, negation, the quantifiers and so on. (We think that the jury is out as to whether logical consequence in languages much richer than first order logic is "usefully formalisable.")

5. If meaning has something to do with truth conditions (and it is plausible to think so) then it is clear that the logical consequences involving conjunction, disjunction, negation and so on — explained in terms of the *cases* in which claims of these forms are true — will have a connection with the meanings of claims of these forms.

6. A lot of information about the truth of claims in different cases is indeed available in principle *a priori* (that a conjunction is true in a situation if and only if its conjuncts are, for example), so it follows that a lot of logical consequence *can* be studied *a priori*.

7. And finally, since $A \wedge B$ is (in general) true in different cases to $A \vee B$, *however* you give an account of cases, their logical consequences will differ. The same is true for the other examples Nolan picked.

So, this sample of intuitions about consequence is indeed *explained* by an appeal to (V). While we agree that (V) alone might not capture *everything* about consequence, it is clearly a central notion. We argue that it is also a notion which may be made precise in a number of equally acceptable ways.

But Consequence Just *Isn't* Relative

Objection: "Intuitively, the notion of consequence just is not relative. According to pluralism, however, consequence is very much a relative notion — relative to specifications of (V)'s cases, *et cetera*. Accordingly, pluralism is counterintuitive."[26]

Reply: Pluralism maintains that the intuitive notion is given entirely in (V). Pluralism maintains, further, that (V) is *very* neutral with respect to its cases — that "pretheoretic intuition" says precious little about what (V)'s cases amount to. But at this stage the (current) objector disagrees. According to the running objection, (V) is far from neutral with respect to its cases; indeed, (V) constrains the field down to one. What can we say?

Ultimately, the case for pluralism may well come down to whether it makes the most sense of contemporary Logic — or at least those branches of Logic concerned with *entailment*. As above, we think that it does, but we will not belabour the point here. Moreover, we've already noted that, according to pluralism, (V) exhausts the pretheoretic notion of consequence, and (V) is neutral with respect to its cases. Beyond what we have said in earlier sections, is there any further "intuitive support" we can give for such claims?

Maybe just this. All sides will agree, we think, that, intuitively, Logic is supposed to account for *good argument* — that is, Logic is the study of "good

[26] Thanks to Ed Zalta for discussion on this point.

arguments", where here we restrict our attention to deductively good arguments. In Logic, we call good arguments 'valid'. In the end, then, Logic tells us that an argument is a *good one* iff there's no case in which its premises are true but its conclusion fails to be true. But, of course, there certainly seem to be different senses of 'good argument', even when this is restricted to deductive cases. Intuitively, the argument from *a is red* to *a is coloured* is a good argument. On the other hand, there's also an equally intuitive sense in which the argument *is not* a good one. In effect, pluralism says that intuition stops there, and that intuition about whether the argument is *really* valid *ought* to stop there. There is a sense in which the colour argument is good, and another sense in which it is not good. Why should we think that there is one absolute sense of validity to use to evaluate reasoning?

Fundamentally, we respond to this objection by *resisting* it. We think that there is a sense in which the argument from $A \vee B$ and $\sim A$ to B is valid, and another sense in which it is not. To ask whether it is *really* valid is beside the point. If the objector can explain *why* it is valid (or why it is not) by spelling out a stronger account of validity which will decide the matter, we will simply point the objector to our *alternative* accounts which are also explications of the fundamental notion of logical consequence.

Is This Conceptual Analysis?

Question: "What is the status of your investigation? Are you engaging in a *conceptual analysis* of the concept of logical consequence?"

Reply: We are not as clear on what counts as conceptual analysis as we would like, and for this reason our remarks on this question are tentative. But insofar as Tarski was doing conceptual analysis in his "On The Concept of Logical Consequence", we suppose that we are to — sort of, and if so, slightly better than Tarski did (or so we hope).[27]

As many have noted, Tarski's aim was in large part to give an account, an "analysis," of the "intuitive" or "ordinary" notion of consequence. Etchemendy repeats the story well:

> Tarski begins his article by emphasizing the importance of the intuitive notion of consequence to the discipline of logic. He dryly notes that the introduction of this concept into the field "was not a matter of arbitrary decision on the part of this or that investigator" (1956, p. 409). The point is that when we give a precise account of this notion, we are not arbitrarily defining a new concept whose properties we then set out to study — as we are when we introduce, say, the concept of a group, or that of a real closed field. It is for this reason that Tarski takes as his goal an account of consequence that remains faithful to the ordinary, intuitive concept from which we borrow the name. It is for this reason that the task becomes, in large part, one of conceptual analysis. [17, page 2]

[27] Of course, we take Tarski's own "analysis" to be captured more or less in (V). The apparent difference between him and us is that we, unlike him, take (V) to be entirely neutral — or, at least, *very* neutral — with respect to (V)'s cases. (Note: This is an *apparent* difference; historians may differ with our reading.)

There is no doubt that we are *not* introducing a new concept and recommending that people study that — at least not in the way that you might be when you introduce the notion of a real closed field.

What we *are* saying is that (V) captures all that is given in the *pretheoretic* "intuitive" notion to which Tarski held his own account accountable. Somewhere in here, we're probably doing as much conceptual analysis as Tarski was doing — and we're probably as clear about the meaning of 'conceptual analysis' as Tarski himself. Where we take a stand is in saying that (V), though an important guide (*the* guide) in Logic, does not constrain the field down to one candidate, but rather, leaves the field open for a great deal of "play".

7 Conclusion

Logic is a matter of truth preservation in all cases. Different *logics* are given by different explications of the cases in question. This account of the nature of logical consequence sheds light on debates between different formal logics. They arise from different accounts of the "cases" in which claims are true or not. Once this realisation is made, apparent disagreements between some formal logics are shown to be just that: merely *apparent*. A number of different formal logics (in particular, classical logics, relevant logics and intuitionistic logics) each have their place in formalising and regulating inference. Each is an elucidation of our pretheoretic, intuitive notion of logical consequence.[28]

References

[1] Jon Barwise and John Etchemendy. "Information, Infons, and Inference". In Robin Cooper, Kuniaki Mukai, and John Perry, editors, *Situation Theory and its Applications 1*, number 22 in csli Lecture Notes, pages 33–78. csli Publications, Stanford, 1990.

[2] Jon Barwise and John Perry. *Situations and Attitudes*. MIT Press, Bradford Books, 1983.

[3] Jon Barwise and John Perry. "Shifting Situations and Shaken Attitudes". *Linguistics and Philosophy*, 8:105–161, 1985.

[4] Nuel Belnap and Mitchell Green. "Indeterminism and the Thin Red Line". *Philosophical Perspectives (Logic and Language)*, 8:365–388, 1994.

[5] Nuel D. Belnap. "A Useful Four-Valued Logic". In J. Michael Dunn and George Epstein, editors, *Modern Uses of Multiple-Valued Logics*, pages 8–37. D. Reidel, Dordrecht, 1977.

[6] Jonathan Bennett. "Entailment". *The Philosophical Review*, 68:197–236, 1969.

[7] Errett Bishop. *Foundations of Constructive Analysis*. McGraw-Hill, 1967. Out of print. A revised and extended version of this volume has appeared [8].

[8] Errett Bishop and Douglas Bridges. *Constructive Analysis*. Springer Verlag, 1985.

[9] George Boolos. "Nominalist Platonism". *Philosophical Review*, 94:327–344, 1985.

[10] Douglas S. Bridges. *Constructive Functional Analysis*, volume 28 of *Research notes in mathematics*. Pitman, 1979.

[28] Thanks to James Chase, Daniel Nolan, John Slaney, Ed Zalta, and an audience at the Automated Reasoning Project at the Australian National University for comments on an earlier version of this paper.

[11] Keith Devlin. *Logic and Information.* Cambridge University Press, 1991.

[12] Michael Dummett. *Elements of Intuitionism.* Oxford University Press, Oxford, 1977.

[13] J. Michael Dunn. "Intuitive Semantics for First-Degree Entailments and "Coupled Trees"". *Philosophical Studies*, 29:149–168, 1976.

[14] J. Michael Dunn. "Star and Perp: Two Treatments of Negation". In James E. Tomberlin, editor, *Philosophical Perspectives*, volume 7, pages 331–357. Ridgeview Publishing Company, Atascadero, California, 1994.

[15] J. Michael Dunn. "Generalised Ortho Negation". In Heinrich Wansing, editor, *Negation: A Notion in Focus*, pages 3–26. Walter de Gruyter, Berlin, 1996.

[16] John Etchemendy. "Tarski on truth and logical consequence". *Journal of Symbolic Logic*, 53(1):51–79, 1988.

[17] John Etchemendy. *The Concept of Logical Consequence.* Harvard University Press, Cambridge, Mass., 1990.

[18] John F. Fox. "Truthmaker". *Australasian Journal of Philosophy*, 65:188–207, 1987.

[19] William H. Hanson. "The Concept of Logical Consequence". *The Philosophical Review*, 106:365–409, 1997.

[20] Richard C. Jeffrey. *Formal Logic: its scope and its limits.* McGraw Hill, Third edition, 1991.

[21] G. Kreisel. "Informal Rigour and Completeness Proofs". In Imre Lakatos, editor, *Problems in the philosophy of mathematics.* North-Holland, Amsterdam, 1967.

[22] Saul A. Kripke. *Naming and Necessity.* Harvard University Press, Cambridge, MA, 1972.

[23] Mark Lance. "Two Concepts of Entailment". *Journal of Philosophical Research*, 20:113–137, 1995.

[24] David K. Lewis. *On the Plurality of Worlds.* Blackwell, Oxford, 1986.

[25] Michael J. Loux, editor. *The Possible and the Actual: Readings in the Metaphysics of Modality.* Cornell University Press, 1979.

[26] Jan Łukasiewicz. "On Determinism". In L. Borkowski, editor, *Selected works.* North Holland, Amsterdam, 1970.

[27] E. P. Martin and R. K. Meyer. "Solution to the P-W problem". *Journal of Symbolic Logic*, 47:869–886, 1982.

[28] Vann McGee. "Two Problems with Tarski's Theory of Consequence". *Proceedings of the Aristotelian Society*, 92:273–292, 1992. 1992.

[29] Robert K. Meyer and J. Michael Dunn. "E, R and γ". *Journal of Symbolic Logic*, 34:460–474, 1969.

[30] Robert K. Meyer and Errol P. Martin. "Logic on the Australian Plan". *Journal of Philosophical Logic*, 15:305–332, 1986.

[31] Robert K. Meyer and Errol P. Martin. "On Establishing the Converse". *Logique et Analyse*, 139–140:207–222, 1992.

[32] Ray Mines, Fred Richman, and Wim Ruitenburg. *A Course in Constructive Algebra.* Springer-Verlag, 1988.

[33] Kevin Mulligan, Peter Simons, and Barry Smith. "Truth-Makers". *Philosophy and Phenomenological Research*, 44:287–321, 1984.

[34] D. Nelson. "Constructible Falsity". *Journal of Symbolic Logic*, 14:16–26, 1949.

[35] W. H. Newton-Smith. *Logic: an introductory course*. Routledge and Kegan Paul, 1985.

[36] Greg O'Hair. "Logical Consequence and Model-Theoretic Consequence". *Logique et Analyse*, 139–140:239–249, 1992.

[37] Graham Priest. "The Logic of Paradox". *Journal of Philosophical Logic*, 8:219–241, 1979.

[38] Graham Priest. *In Contradiction: A Study of the Transconsistent*. Martinus Nijhoff, The Hague, 1987.

[39] Graham Priest. "Logic, Nonstandard". In Donald M. Borchert, editor, *Encyclopedia of Philosophy, Supplement*. Macmillan Reference, 1996.

[40] Stephen Read. "What is wrong with Disjunctive Syllogism?". *Analysis*, 41:66–70, 1981.

[41] Stephen Read. "Formal and Material Consequence". *Journal of Philosophical Logic*, 23:247–265, 1994.

[42] Stephen Read. *Thinking about Logic*. Oxford University Press, 1995.

[43] Greg Restall. "Łukasiewicz, Supervaluations and the Future". Presented at the 1996 Łukasiewicz in Dublin Conference. To appear in *Łukasiewicz in Dublin*, Oxford University Press.

[44] Greg Restall. "A Note on Naïve Set Theory in **LP**". *Notre Dame Journal of Formal Logic*, 33:422–432, 1992.

[45] Greg Restall. "Deviant Logic and the Paradoxes of Self Reference". *Philosophical Studies*, 70:279–303, 1993.

[46] Greg Restall. "Four Valued Semantics for Relevant Logics (and some of their rivals)". *Journal of Philosophical Logic*, 24:139–160, 1995.

[47] Greg Restall. "Truthmakers, Entailment and Necessity". *Australasian Journal of Philosophy*, 74:331–340, 1996.

[48] Greg Restall. "Negation in Relevant Logics: How I stopped worrying and learned to love the Routley star". In Dov Gabbay and Heinrich Wansing, editors, *What is Negation?* To be published by Kluwer, 1998.

[49] Greg Restall. "Ways Things Can't Be". *Notre Dame Journal of Formal Logic*, 39:?–?, 1998. To Appear.

[50] Fred Richman. "Interview with a constructive mathematician". *Modern Logic*, 6:247–271, 1996.

[51] Richard Routley. "Relevantism, Material Detachment, and the Disjunctive Syllogism Argument". *Canadian Journal of Philosophy*, 14:167–188, 1984.

[52] Richard Routley and Valerie Routley. "Semantics of First-Degree Entailment". *Noûs*, 3:335–359, 1972.

[53] Manuel García-Carpintero Sánchez-Miguel. "The Grounds for the Model-theoretic Account of the Logical Properties". *Notre Dame Journal of Formal Logic*, 34:107–131, 1993.

[54] Gila Sher. *The Bounds of Logic*. MIT Press, 1991.

[55] Robert Stalnaker. *Inquiry*. Bradford Books, MIT Press, 1984.

[56] W. W. Tait. "Against Intuitionism: Constructive Mathematics is Part of Classical Mathematics". *Journal of Philosophical Logic*, 12:173–195, 1983.

[57] W. W. Tait. "Truth and Proof: The Platonism of Mathematics". *Synthese*, 69:314–370, 1986.

[58] Alfred Tarski. *Logic, Semantics, Metamathematics: papers from 1923 to 1938*. Clarendon Press, Oxford, 1956. Translated by J. H. Woodger.

[59] Neil Tennant. "The Transmission of Truth and the Transitivity of Deduction". In Dov Gabbay, editor, *What is a Logical System?*, volume 4 of *Studies in Logic and Computation*, pages 161–177. Oxford University Press, Oxford, 1994.

[60] Richmond Thomason. "Indeterminist Time and Truth-Value Gaps". *Theoria*, 3:264–281, 1970.

[61] Ludwig Wittgenstein. *Remarks on the Foundations of Mathematics*. MIT Press, 1967. Edited by G. H. von Wright, R. Rees and G. E. M. Anscome.

[62] Edward N. Zalta. "Twenty-Five Basic Theorems in Situation and World Theory". *Journal of Philosophical Logic*, 1993:385–428, 1993.

[63] Edward N. Zalta. "A Classically-Based Theory of Impossible Worlds". *Notre Dame Journal of Formal Logic*, 39:?–?, 1998. To Appear.

JC Beall
School of Philosophy, Faculty of Arts, The University of Tasmania,
GPO Box 252-41, Hobart, Tasmania 7001, Australia.
j.c.beall@utas.edu.au
http://www.utas.edu.au/docs/humsoc/philosophy/JCBeall.html

Greg Restall
Department of Philosophy, Macquarie University,
NSW 2109, Australia.
Greg.Restall@mq.edu.au
http://www.mq.edu.au/~phildept/staff/grestall/