Nuel D. Belnap, Jr.
and
Gerald J. Massey

Semantic Holism

Abstract. A bivalent valuation is *snt* iff sound (standard *PC* inference rules take truths only into truths) and non-trivial (not all wffs are assigned the same truth value). Such a valuation is *normal* iff classically correct for each connective. Carnap knew that there were non-normal snt valuations of *PC*, and that the gap they revealed between syntax and semantics could be "jumped" as follows. Let VAL_{snt} be the set of snt valuations, and VAL_{nrm} be the set of normal ones. The bottom row in the table for the wedge '$\vee$' is not semantically determined by VAL_{snt}, but if one deletes from VAL_{snt} all those valuations that are not classically correct at the aforementioned row, one jumps straights to VAL_{nrm} and thus to classical semantics. The conjecture we call *semantic holism* claims that the same thing happens for any semantic indeterminacy in any row in the table of any connective of *PC*, i.e., to remove it is to jump straight to classical semantics. We show (i) why semantic holism is plausible and (ii) why it is nevertheless false. And (iii) we pose a series of questions concerning the number of possible steps or jumps between the indeterminate semantics given by VAL_{snt} and classical semantics given by VAL_{nrm}.

1. Historical Context

Rudolf Carnap's influential 1942 book *Introduction to Semantics* has so fascinated logicians that, for the most part, they have overlooked or forgotten its provenance. According to his own account, Carnap wrote it to set the stage for the ideas and results to be found in his 1943 book *Formalization of Logic.* Though published later, the 1943 book was drafted earlier than the 1942 one, in the autumn of 1938.

What were the ideas and results that Carnap deemed so important that he composed another substantial book to pave the way for them? In a word, *misalignment.* Carnap had discovered that classical syntax and classical semantics don't line up quite so neatly as nearly everyone seemed to think. For example, the consistency and completeness metatheorems for classical propositional logic (*PC*) were taken uncritically, even unconsciously, to entail that *sound* (in the sense that the inference rules of *PC* permit only truths to be deduced from truths) *bivalent valuations* (functions that map the wffs into the two truth values) of the wffs of *PC* must conform every-where to the classical truth tables for the connectives. (Throughout this paper we will be concerned only with *bivalent* valuations, so we will hence-forth make bivalence part of the concept of a *valuation.*)

By a *normal* valuation Carnap meant a sound valuation in which all sentential connectives conform everywhere to the prescriptions of their classical truth tables. By a *non-normal* valuation Carnap meant a sound valuation in which at least one sentence connective deviates in at least one instance from the prescriptions of its classical truth table. (Hereafter, we will use the terms *normal valuation* and *non-normal valuations* in the foregoing Carnapian sense. Note in particular that both normal and non-normal valuations are sound.) What had profoundly unsettled Carnap was his discovery of *non-normal valuations*, i.e., sound valuations that here and there violate the prescriptions of classical truth tables.

Carnap's non-normal valuations come in two varieties. Those of the first kind violate the semantical law of non-contradiction, which requires of a sentence and its negation that at least one of them be false. There happens to be only one non-normal (sound) valuation of this type, to wit: the function that assigns truth to every wff. This non-normal valuation renders all connectives truth functional but in a bland way, for it simply turns them all into constant-truth operators. But however trivial this particular non-normal valuation might be in itself, it posed for Carnap the deeply arresting problem of how to exclude it by syntactical means alone.[1]

Non-normal valuations of the second sort violate the semantical law of excluded middle, which demands of a sentence and its negation that at least one be true. The existence of such non-normal valuations is not obvious, although it turns out that their number is legion. As an example of a non-normal valuation of this latter variety, we mention a particularly important one that was given already by Carnap: the valuation V^*, where V^* maps a wff B onto t or f according as B is or is not a theorem of *PC*.

Since only theorems are deducible from theorems, V^* is obviously sound. That V^* does not conform everywhere to classical truth tables is equally evident. V^* assigns f to 'p' but also to '$\sim p$' because neither of these wffs is a *PC* theorem, so the semantics of negation proves to be abnormal in V^*. The semantics of disjunction is warped by V^* too, because V^* assigns t to the theorem '$p \vee \sim p$' while assigning f to both disjuncts. The semantics of the conditional and the biconditional connectives fare no better. Though its antecedent and consequent are both f in V^*, the conditional '$p \supset q$' comes out f in V^* too. Though both of its sides agree in truth value (both are false), in V^*, the biconditional '$p \equiv q$' also comes out f in V^*. Nor does exclusive disjunction escape semantically unscathed. For example, the exclusive disjunction '$p \not\equiv \sim p$' of 'p' and '$\sim p$' is t in V^* although both disjuncts come out f in V^*. The semantics of the two binary Sheffer connectives, non-conjunction and non-disjunction, fall victim to V^* too. For example, both '$p|q$' and '$p \downarrow q$' come out false in V^*, although both sides of both formulas are false in V^*.

Of the familiar classical connectives, only conjunction is stalwart enough to fend off non-normal valuations of both Carnapian varieties. Notice, now, that

[1] Carnap was interested in harmonizing semantics and syntax for quantifiers as well as for sentential connectives. We mention this now because we do not further touch on this aspect of his enterprise.

with respect to non-normal valuations, conjunction and disjunction fare differentially. This fact surprised Carnap, even shocked him. Syntax ought to be so wrought, he believed, that non-normal valuations cannot arise. If it were so crafted, these two connectives would share the same semantical fate in view of their perfect semantical symmetry from a classical perspective.[2]

All the results so far mentioned can be found in Carnap's 1943 book. Some of them have been rediscovered in recent years by several people, most notably by James McCawley whose work has rekindled interest in the question of the degree to which classical syntax and classical semantics line up. In his book *Everything that Linguists have Always Wanted to Know about Logic*, McCawley poses what is essentially Carnap's problem this way. To what degree does insistence on the soundness (relative to the inference rules of *PC*) of valuations force the classical semantical interpretation of the sentence connectives on us? McCawley answers this question in considerable detail, though all his results can be found already in Carnap.

But as Alonzo Church pointed out in his 1944 *JSL* review of Carnap's 1943 book, Carnap also did not break virgin ground. Church reminded logicians that such contributors to the algebra of logic as E. V. Huntington and B. A. Bernstein had already in the 1920s and early 1930s studied non-normal valuations of *PC* and had, like Carnap later, made proposals for eliminating them.[3]

It is in fact only a short step from Boolean algebras to non-normal valuations of *PC*. Drop the anomalous Boolean postulate that says the algebra must contain at least two elements and you admit the one-element "Boolean algebra" that induces immediately Carnap's solitary non-normal valuation of the first variety. Take any Boolean algebra that boasts of more than two elements and you can construct from it a non-normal valuation of Carnap's second variety as follows. Map the variables of *PC* into the elements of the algebra (Boolean elements) in whatever way you like so long as at least one element other than the Boolean zero or the Boolean unit has a variable mapped onto it. Then, map the complex wffs into the Boolean elements thus. Map a disjunction (conjunction) onto the Boolean join (Boolean meet) of the Boolean elements mapped with its disjuncts (conjuncts). Map a negation onto the Boolean complement of the Boolean element mapped with the negated subformula. Finally, produce the desired non-normal valuation by assigning truth to all those wffs mapped onto the Boolean unit element, and by assigning falsehood to all the rest.

---

[2] In a probing 1944 review, Church lodged objections to the new syntactic structures of Carnap's *Formalization of Logic*, maintainig that "they are foreign to elementary syntax, and may not be used in the construction of a calculus" (p. 498). We do not plan to comment on Carnap's innovations, so we here indicate only that the interested reader will find Carnap's suggestions defended and furthered in the 1978 book *Multiple Conclusion Logic* by D. J. Shoesmith and T. J. Smiley, as well as in works by others cited by them, e.g., by W. Kneale and by D. Scott.

[3] Copious citations to this literature are to be found in Church's review.

In his review of Carnap's 1943 book, Church exploited the aforementioned connection between Boolean algebras and non-normal valuations to generate a non-normal (sound) interpretation of *PC* that is more concrete than Carnap's abstract non-normal valuations. Church started with the fourmember Boolean algebra whose elements are the numbers 7, 14, 21, and 42, and in which the Boolean join (Boolean meet) of two elements is their gcd (l cm), and in which the Boolean complement of an element is the number obtained by dividing 294 by the element. In this algebra, 7 is the Boolean unit, so the non-normal interpretation obtained by the above-described procedure assigns truth to all and only those wffs that ultimately get mapped onto 7. For example, '$p \vee \sim p$' is such a wff; no matter which of the four elements 'p' was originally mapped onto, the formula '$p \vee \sim p$' will be mapped onto the Boolean unit 7. Now Church "concretized" the example thus. If the wff B is mapped onto the Boolean element $\boldsymbol{n}$, Church stipulated that B is to express the proposition that there are $\boldsymbol{n}$ days in a week. For example, let 'p' be mapped originally onto 14. Then, '$\sim p$' will be mapped onto 21, and '$p \vee \sim p$' will be mapped onto 7. So, both 'p' and '$\sim p$' express false propositions, viz., that there are 14 and 21 days, respectively, in a week, whereas the disjunction of these two false formulas expresses the true proposition that there are 7 days in a week.

2. The Plausibility of Semantic Holism

We set forth here, by means of semantic tables, what Carnap already knew, or at least probably already knew, about the semantic behavior of the singulary and binary sentence connectives of PC relative to sound valuations.[4] From this point on, we shall purge semantics of the solitary non-normal valuation of Carnap's first variety, the trivial valuation that maps every wff onto truth, by restricting our attention to *non-trivial* valuations, i.e., valuations that map the wffs of *PC onto* the truth values t and f. (Reminder: all valuations are understood to be bivalent, and all normal and non-normal valuations are understood to be sound relative to the classical inference rules of *PC*.) Here, then, are the semantic tables, followed by an explanation of how to interpret them (a question mark means only that there is no standard symbol for the connective with the given table):

A	$?A$	$?A$	$\sim A$	$?A$
t	t	t	f	f
f	t	f	(t)	f

A	B	$A?B$	$A \vee B$
t	t	t	t
t	f	t	t
f	t	t	t
f	f	t	(f)

[4] What we mean by "probably knew" is that the few facts depicted in the semantic tables of the present paper that are not to be found explicitly in Carnap are so close to the surface that one can tell from what Carnap did say that he would not have overlooked them.

$A \subset B$	$A?B$	$A \supset B$	$A?B$	$A \equiv B$	$A \& B$	$A \mid B$
t	t	t	t	t	t	f
t	t	f	f	f	f	(t)
f	f	t	t	f	f	(t)
(t)	f	(t)	f	(t)	f	(t)

$A \not\equiv B$	$A?B$	$A \not\supset B$	$A?B$	$A \not\subset B$	$A \downarrow B$	$A?B$
f	f	f	f	f	f	f
(t)	(t)	(t)	f	f	f	f
(t)	f	f	(t)	(t)	f	f
(f)	(t)	f	(t)	f	(t)	f

To help the reader interpret the semantic tables, we enter a few definitions. First, by a (value-assignment) *row* of a table, we will mean an ordered pair $\langle \#; r \rangle$, where $\#$ is an n-ary connective and r is a sequence $v_1, \ldots, v_n$ of truth values. When context makes it clear which connective is meant, we will sometimes speak of r (the sequence of truth values) as if it were itself the row. Second, we say that a valuation V is *truth functional at a row* $\langle \#; v_1, \ldots, v_n \rangle$ iff, for arbitrary wffs $A_1, \ldots, A_n$, $B_1, \ldots, B_n$, the valuation V assigns the same value to '$\#(A_1, \ldots, A_n)$' and to '$\#(B_1, \ldots, B_n)$' whenever, for each i $(1 \leqslant i \leqslant n)$, the valuation V assigns the value v_i to both A_i and B_i. Third, we say that a valuation V is *classically correct at a row* $\langle \#; v_1, \ldots, v_n \rangle$ iff, for every sequence of wffs $A_1, \ldots, A_n$ to which V respectively assigns the values $v_1, \ldots, v_n$, the value of '$\#(A_1, \ldots, A_n)$' in V is the classical value v (i.e., v is the value assigned to '$\#(A_1, \ldots, A_n)$' by the classical truth table for the connective $\#$ when $A_1, \ldots, A_n$ have been assigned the values $v_1, \ldots, v_n$ respectively). Finally, we say that a set *VAL* of valuations *semantically determines a row* $\langle \#; v_1, \ldots, v_n \rangle$ iff every valuation in *VAL* is classically correct at $\langle \#; v_1, \ldots, v_n \rangle$.

We pause to respond to a possible objection to our choice of the *definiendum* term 'semantically determines'. Someone might think that the term is ill-chosen on the grounds that there may be a set *VAL** of sound valuations that intuitively "determines" a tabular row $v_1, \ldots, v_n$ for an n-ary connective $\#$, but does so *anti-classically*, giving always the wrong value (the value opposite to the classical value) to '$\#(A_1, \ldots, A_n)$' when it assigns the values $v_1, \ldots, v_n$ to $A_1, \ldots, A_n$ respectively. But in fact there can be no such set *VAL** of sound valuations, and so the question whether a sound valuation is classically correct at a row is equivalent to whether it is truth functional at the row. Hence, a set of sound valuations semantically determines a tabular row in our technical sense just in case it semantically determines the row in the intuitive sense of this phrase.

PROOF of the foregoing claim: Let $\langle @; v_1, \ldots, v_n \rangle$ be a tabulator row, and let v be the classical value of the wff '$@(A_1, \ldots, A_n)$' when $A_1, \ldots, A_n$ are assigned the values $v_1, \ldots, v_n$ respectively. Let V be any sound (nontrivial) valuation. Let T be the PC theorem '$p \vee \sim p$', let F be the PC antitheorem '$p \& \sim p$', and let C_i be T or F according as v_i is t or f. Then, the wff '$@(C_1, \ldots, C_n)$' has the classical value v in the valuation V, while its argument formulas $C_1, \ldots, C_n$ are assigned the values $v_1, \ldots, v_n$ respectively by V. Hence, no sound valuation can determine a tabular row anti-classically.

Armed now with the foregoing definitions, we return to our explanation of the semantic tables above. An *unadorned* outcome entry in a table, i.e., an outcome entry not flanked by parentheses, indicates a tabular row that is semantically determined by the set of all sound valuations. For future ease of reference, we will let VAL_{snt} be the set of all sound (non-trivial) valuations. We also let VAL_{nrm} be the set of all normal valuations, which of course are just those sound (non-trivial) valuations that semantically determine every row of every connective. For example, the unadorned outcome entries on the first three rows in the semantic table for the wedge indicate that, in any sound valuation, whether normal or non-normal, a disjunction will be true if even one of its disjuncts is true.

By contrast, a *parenthesized* outcome entry on a row of a semantic table indicates a tabular row that is not semantically determined by VAL_{snt}. For example, the entry '(f)' on the bottom row of the table for the wedge signifies that there is a sound valuation of PC relative to which some disjunction with two false disjuncts comes out abnormally, i.e., comes out true in the valuation. To see that this is so, note that the special sound valuation V^* assigns t to the disjunction '$p \vee \sim p$' although it assigns f to its two disjuncts 'p' and '$\sim p$'. Of course, it follows from what we proved above that V^*, like every sound valuation, will sometimes behave classically at this row. That is, there must be some disjunctions that are themselves false in V^* and both of whose disjuncts are also false in V^*.

Similar remarks hold for the parenthesized outcome entry '(t)' on the second row of the table for exclusive disjunction. This row is not semantically determined by VAL_{snt} but nevertheless every sound valuation behaves properly (i.e., classically) at this row with respect to some exclusive disjunctions whose left and right disjuncts are respectively true and false in the valuation.

It is child's play to establish the facts depicted in the semantic tables above. For example, the unadorned outcome entries on the first three rows of the table for the wedge signify that VAL_{snt} sementically determines these rows. To see that this is true, one need only remember that in PC a disjunction can be deduced from either of its disjuncts. Hence, a valuation that assigned t to even one of the disjuncts while assigning f to a disjunction would fail to be sound, for one could then deduce a falsehood from a truth. Moreover, to justify the parenthesized outcome entries in the various semantic tables, i.e., to show that VAL_{snt} does not semantically determine these rows, one need appeal to no non-normal valuation beyond V^* to illustrate the requisite semantic abnormalities.

We will say that a tabular row r is a *semantic indeterminacy* iff r is *not* semantically determined by VAL_{snt}. For example, we saw above that the row $\langle$'$\vee$'; $f, f\rangle$, the bottom row of the table for the wedge, is a semantic indeterminacy because it is not semantically determined by VAL_{snt}. *Relative to a set VAL of sound valuations, the result of removing a semantic indeterminacy r,* where $r = \langle \#, v_1, \ldots, v_n\rangle$, will be understood to be the set $VAL_\#$ of valuations that remain when one has ejected from VAL all and only those valuations that fail to be classically correct at the row r. (In other words, $VAL_\#$ is the set of all those valuations in VAL that are classically correct at r.)

Carnap had already noted that, if you start with VAL_{snt} and then remove the semantic indeterminacy in negation, all semantic indeterminacies in all connectives disappear. That is, Carnap saw that VAL_{neg}, the set of all sound valuations that are classically correct at the bottom row of the table for negation, is just VAL_{nrm}. (In view of the fact that conjunction exhibits no semantic indeterminacies, the foregoing observation is but a corollary of the definitional completeness in PC of negation and conjunction.) That is to say, if you require of valuations not only that they be sound but also that they assign t to the negation of any formula to which they have assigned f, you will be left with only normal valuations, the valuations in which the semantics of every sentence connective conforms to its classical truth table.

Further, Carnap recognized that if you remove the bottom-row semantical indeterminacy in any of the familiar binary connectives, you thereby cause all semantic indeterminacies in all connectives to disappear and so obtain classical semantics. This result is somewhat surprising in the cases of connectives like the wedge that do not combine with the ampersand to form a definitionally complete set of connectives, but it is nonetheless easy to prove. For example, suppose we have removed the semantic indeterminacy in the wedge by discarding from VAL_{snt} all valuations in which the disjunction of a pair of false formulas is not invariably false. Let VAL_v be the set of valuations that remain, i.e., the members of VAL_v are those sound (non-trivial) valuations (members of VAL_{snt}) in which the disjunction of any pair of false formulas is itself false. We now show that there remains no semantic indeterminacy in negation, and *a fortiori* in no other connective, relative to VAL_v by showing that an arbitrary member V of VAL_v must assign t to the negation of any formula to which it has assigned f. For the purpose of indirect proof, suppose the opposite is the case, i.e., suppose that V assigns f to at least one wff '$\sim A$' when it has also assigned f to the subformula A. All sound valuations must assign t to all theorems of PC, so V assigns t to '$A \vee \sim A$'. But *V must assign f* to this disjunction because it assigns f to both of its disjuncts, and so we have our contradiction.

One of us (GJM) has appealed to these results to make a case for the classical interpretation of the sentence connectives of PC.[5] The argument

[5] See section v. of Massey's "The Pedagogy of Logic: Humanistic Dimensions".

goes like this. If you view the inference rules of *PC* as sound, i.e., as never leading from truths to falsehood, and if you insist on a bivalent (non-trivial) interpretation in which the negation of a false sentence is invariably true, you have no choice but to accept the classical truth tables for the sentence connectives. Classical semantics is similarly thrust upon you if you insist that disjunctions be false when both disjuncts are false, or if you insist that biconditionals be true whenever both sides are false, and so on for all the bottom-row indeterminacies in the semantic tables above.

The bottom-row semantic indeterminacies in the above tables are not the only semantic indeterminacies whose removal yields classical semantics. It is easy to show of *each* of the semantic indeterminacies in these tables that its removal eliminates all semantic indeterminacies, i.e., that its removal yields classical semantics. But these tables are limited to singulary and binary connectives. Many connectives of higher degree exhibit semantic indeterminacies as well. Indeed, for every $n > 0$, there are n-ary connectives with as many as $2^n - 1$ semantic indeterminacies in their tables [See result (i) below], so there is not even an upper bound on the number of semantic indeterminacies a single connective can exhibit. Will removing any semantic indeterminacy in any connective, therefore, eliminate all semantic indeterminacies in all connectives? Is classical semantics thus *holistic*?

Let us call a semantic indeterminacy r in a connective $\#$ *systemic* if and only if its removal causes all semantic indeterminacies in all connectives to disappear, i.e., if and only if every sound (non-trivial) valuation that is classically correct for $\#$ at row r is normal. Then the fact that every semantic indeterminacy in every singulary and binary connective is systemic by itself suggests that all semantic indeterminacies in all connectives are systemic too. This suggestion, that classical semantics is unremittingly holistic, is reinforced by a number of results, some of which we now set forth:

(i) *The table for each n-place* ($n > 1$) *non-conjunction connective* (*the n-place analog of Sheffer's stroke*) *exhibits a semantic indeterminacy on every row except the top one.* (*So that reference to row numbers are unambiguous, we assume that value-assignment entries in tables are made lexicographically, with 't' taking precedence over 'f'.*) *Furthermore, each of these semantic indeterminacies is systemic.*

PROOF. Let the n-place non-conjunction connective be '$|^n$', i.e., let '$|^n(A_1, \ldots, A_n)$' be classically equivalent to '$\sim(A_1 \& \ldots \& A_n)$'. Let $v_1, \ldots, v_n$ be any tabular row r for '$|^n$' other than the top row. Let T be some particular *PC* theorem, say '$p \vee \sim p$'. Let V be any (sound) valuation in which some wff A and its negation '$\sim A$' are both false (we have already seen that there are such valuations). Let A_i^* ($1 \leqslant i \leqslant n$) be T or A according as v_i is t or f. Clearly, '$|^n(A_1^*, \ldots, A_n^*)$' and '$\sim A$' are interdeducible in *PC* and so agree in value in any sound valuation. Hence, both of these wffs come out f in V. But, in the classical truth table for '$|^n$', the outcome entry on row r is t, so the classical value of '$|^n(A_1^*, \ldots, A_n^*)$' is t on row r. Therefore, there is a semantic indeterminacy at row r for '$|^n$'. Moreover, the *PC* interdeducibility of '$\sim A$' and '$|^n(A_1^*, \ldots, A_n^*)$' makes it obvious that removing the semantic indeterminacy at row r for '$|^n$' eliminates the semantic indeterminacy in negation, and so eliminates all semantic indeterminacies.

(ii) *The removal of a semantic indeterminacy in an n-ary connective* $\otimes$, $n > 1$, *at any non-bottom row r eliminates a semantic indeterminacy in some connective of smaller degree.*

PROOF. Let 'A_1', ..., 'A_n' be the schematic argument formulas in the classical truth table for $\otimes$. At least one of these formulas, say A_j, will be assigned t on the value-assignment row r, because the only row where all schematic wffs are assigned f is the bottom one. Let row r be $v_1, \ldots, v_n$. Let $w_1, \ldots, w_k$ be, in top-to-bottom order, the outcome entries for '$\otimes(A_1, \ldots, A_n)$' on the k rows of the classical truth table for $\otimes$ on which A_j is assigned the value t (k will be equal to $2^n/2$). Then, the biconditional of '$\otimes(A_1, \ldots, A_n)$' and '$@(A_1, \ldots, A_{j-1}, A_{j+1}, \ldots, A_n)$' is *PC* deducible from A_j, where @ is the $(n-1)$-place connective in whose classical truth table the outcome column is given (from top to bottom) by $w_1, \ldots, w_k$. Hence, there must be a semantic indeterminacy on the row r^* of the semantic table for @, where r^* is $v_1, \ldots, v_{j-1}, v_{j+1}, \ldots, v_n$, because when A_j is assigned the value t the two sides of the aforementioned biconditional must agree in truth value on row r^*. Now, substituting the *PC* theorem T for A_j in the left side of the aforementioned biconditional turns that biconditional into a *PC* theorem. Thus, inspection of the semantic table for the biconditional shows that removal of the semantic indeterminacy in $\otimes$ at row r will eliminate the semantic indeterminacy in @ at row r^*.

(iii) *Let the n-ary connective* #, $n > 1$, *exhibit a semantic indeterminacy at its bottom row, and let the bottom-row entry in the classical truth table for # be t. Then, removal of the bottom-row semantic indeterminacy in # eliminates a semantic indeterminacy in some connective of smaller degree.*

PROOF. We will prove something stronger, viz., that the bottom-row semantic indeterminacy in # is systemic. First, note that there must be an outcome entry f on some row r in the classical truth table for #; otherwise, # would be the n-ary constant-truth connective, which exhibits no semantic indeterminacies. So, let row r be $v_1, \ldots, v_n$. Let F be the *PC* anti-theorem '$p \,\&\, \sim p$'. Let A be an arbitrary wff, and let A_i^* be A or F according as v_i is t or f. Notice, by the way that '$\#(A_1^*, \ldots, A_n^*)$' has been constructed, '$\sim A$' and '$\#(A_1^*, \ldots, A_n^*)$' are *PC* interdeducible since these formulas are classically equivalent. Thus, removing the bottom-row indeterminacy in # also removes the indeterminacy in the tilde, and so the bottom-row semantic indeterminacy in # is systemic. [Note, by the way, that the prrof shows that the row $\langle \#, r \rangle$ is semantically determined by VAL_{snt}.]

(iv) *Let the n-ary connective #, $n > 1$, exhibit a bottom-row semantic indeterminacy, and let f be the outcome value on the bottom row of the classical truth table for #. Then, if the outcome values of even one pair of dual value-assignment rows in the classical truth table for # are both t, removal of the bottom-row semantic indeterminacy in # also eliminates a semantic indeterminacy in some connective of smaller degree.*

PROOF. Again we prove something stronger, viz., that the bottom-row semantic indeterminacy in # is systemic. [To say that two value-assignment rows are *dual* means that each schematic argument formula is assigned opposite truth values on the two rows.] Let $v_1, \ldots, v_n$ be one of the pair of dual rows. Let A be an arbitrary wff, and let A_i^* be A or '$\sim A$' according as v_i is t or f. Then, the formula '$\#(A_1^*, \ldots, A_n^*)$' is a PC theorem and so must receive the value t in *all* sound valuations. Suppose now that we remove the bottom-row semantic indeterminacy in # by discarding all sound valuations that are not classically correct at the bottom row for #. We are left then with the set $VAL_\#$ of sound valuations in which an arbitrary wff of the form '$\#(C_1, \ldots, C_n)$' is assigned f whenever its subformulas $C_1, \ldots, C_n$ are all assigned f. So, if the wff A and its negation '$\sim A$' are both assigned f by a valuation V in $VAL_\#$, the valuation V must also assign f to '$\#(A_1^*, \ldots, A_n^*)$'. From this contradiction it follows that there is no semantic indeterminacy in negation, i.e., that $VAL_\#$ semantically determines the bottom row of the table for negation.

In view of the fact that all semantic indeterminacies in singulary and binary connectives are systemic, the results (i) to (iv), as well as many similar results which we do not bother to reproduce here, make it plausible to suppose that the removal of any semantic indeterminacy in any connective of degree 3 or higher eliminates a semantic indeterminacy in some connective of smaller degree. If this were really so, all semantic indeterminacies would be systemic, and so classical semantics would be holistic in the following very strong sense. Either you must accept all semantic indeterminacies or none of them; if you remove even one of them, you are left with classical semantics, the semantics of classical truth tables. Between accepting all semantic indeterminacies and accepting none, *tertium non datur*. This, in fact, is what one of us (GJM) recently conjectured.

3. The Falsity of Semantic Holism

Plausibility, however strong, falls short of proof. Although one of us (GJM) was unable to prove his semantic-holism conjecture, his conviction that it was true never flagged. But illustrating again the adage that one person's *modus ponens* is another's *modus tollens*, the other one of us (NDB) became convinced that failure to prove the semantic-holism thesis was due not to lack of effort or ingenuity but rather to falsity of the thesis itself. This led him to exploit failure by turning the very difficulties encountered in trying to prove semantic holism into the following counterexample to the thesis:

COUNTEREXAMPLE TO THE THESIS OF SEMANTIC HOLISM. *Let 'v^2' be the ternary connective such that '$v^2(A_1, A_2, A_3)$' comes out true when and only when two or more of the items in the list A_1, A_2, A_3 come out true. Then, there is a semantic indeterminacy on the bottom row of the semantic table for 'v^2'. Moreover, this particular semantic indeterminacy is not systemic, i.e., its removal does not eliminate all semantic indeterminacies. Hence, the thesis of semantic holism is false.*

PROOF. To make it easier to follow the proof, we first display the semantic table for 'v^2', which expresses what might be called *two-out-of-three* disjunction:

A	B	C	$v^2(A, B, C)$
t	t	t	t
t	t	f	t
t	f	t	t
t	f	f	(f)
f	t	t	t
f	t	f	(f)
f	f	t	(f)
f	f	f	(f)

It is easy to verify the four 't' entries in the above table. Nor is it difficult to verify the first three '(f)' entries. For example, to justify the entry '(f)' on the fourth row, note that the biconditional '$v^2(T, B, C) \equiv (B \vee C)$' is a classical tautology and so a *PC*-theorem, where T is some particular *PC* theorem, say '$p \vee \sim p$'. Hence, there is an indeterminacy marked by '(f)' on the fourth row. Similarly, for the '(f)' entries on the sixth and seventh rows. Finally, to justify the bottom entry '(f)' in the table, we need only appeal to the now familiar special non-normal valuation V^*, which assigns t to the *PC* theorem '$v^2(p, p \supset q, p \supset \sim q)$' while simultaneously assigning f to each of its three subformulas 'p', '$p \supset q$', and '$p \supset \sim q$'.

There remains only the task of showing that the bottom-row semantic indeterminacy in 'v^2' is not systemic, i.e., that its removal does not eliminate all semantic indeterminacies in all connectives. All we need to do, clearly, is show that its removal does not eliminate the semantic indeterminacy in negation. To do this, we construct a valuation $V+$ that has the following three properties: (a) $V+$ is a sound valuation; (b) $V+$ assigns f to some wffs *and* to their negations; and (c) $V+$ is classically correct for 'v^2' at its bottom row, i.e., $V+$ assigns f to any two-out-of-three disjunction *whenever* it has assigned f to all three of its disjuncts.

We construct $V+$ as follows. Let V_t be the normal (classical) valuation in which t is assigned to 'p' and f is assigned to every other sentential variable. Let V_f be the normal valuation in which f is assigned to every sentential variable. Let $V+$ assign t to a wff D iff both V_t and V_f assign t to D; otherwise, let $V+$ assign f to D. Clearly $V+$ is a (non-trivial) valuation, i.e., a mapping of the PC wffs *onto* $\{t, f\}$.

To see that the rest of (a) holds, i.e., that the valuation $V+$ is sound, suppose B is PC deducible from $A_1, \ldots, A_k$. Then, if V_t assigns t to all of $A_1, \ldots, A_k$, the valuation V_t must also assign t to B. Similarly, for V_f. Now $V+$ assigns t to all of $A_1, \ldots, A_k$ if and only if both V_t and V_f assign t to all of these formulas. Hence, $V+$ will assign t to B if it assigns t to all of $A_1, \ldots, A_k$.

To see that (b) holds, note that the valuation $V+$ assigns f both to the wff 'p' and to its negation '$\sim p$'. Thus, $V+$ is a non-normal valuation.

Finally, to see that (c) holds, suppose that $V+$ assigns f to the three wffs A, B, and C; we will show that $V+$ also assigns f to the two-out-of-three disjunction '$v^2(A, B, C)$'. Since $V+$ assigns f *to* A, either V_t or V_f must assign f to A; similarly, for the assignments of f to B and to C by $V+$. So, either V_t or V_f must assign f to at least two of the formulas in the list A, B, C. So, either V_t or V_f must assign f to the two-out-of-three disjunction. '$v^2(A, B, C)$'. Therefore, $V+$ must also assign f to this same disjunction. We have shown, therefore, that $V+$ assigns f to a two-out-of-three disjunction whenever it assigns f to each of its three disjuncts, i.e., we have shown that (c) holds. This completes our proof of the counterexample to the thesis of semantic holism.

Let the reader think that only bottom-row indeterminacies can be non-systemic, we observe that there is a non-systemic semantic indeterminacy on the eighth row of the semantic table for the quaternary connective whose classical truth-table outcome column is, from top to bottom.

tttftffftttftfff

which is the sequence of truth values obtained simply by reiterating the classical truth-table outcome column for 'v^2'. Moreover, given any positive integer n, by starting with the classical outcome column for 'v^2' and repeatedly reiterating the resulting column, we can generate the outcome column of the classical truth table for a many-place connective that exhibits more than n non-systemic semantic indeterminacies. There is, in other words, no upper bound on the number of non-systemic semantic indeterminacies that a single sentence connective may exhibit. (We already saw above that there is no upper bound on the number of systemic indeterminacies that a connective may exhibit.)

4. The Problem of Semantic Holism

In this final section we discuss several open questions. To this end, we first advance a few definitions. By *inferential semantics*, let us understand the disconcertingly indeterminate semantics determined by the classical inference rules of PC. Inferential semantics, then, is the semantics generated by VAL_{snt}, which is of course the semantics exhibited in the semantic tables for *PC* sentence connectives, such as those displayed above for the singulary and binary connectives. By *classical semantics* we mean the semantics determined by the classical truth tables for *PC* connectives, which is of course the semantics generated by VAL_{nrm}. By a *semantic path* we mean a finite or denumerable sequence of sets of sound valuations the first member of which is VAL_{snt} and in which every pair VAL_i and VAL_{i+1} of adjacent members are related as follows: there is a row r of some connective $\#$ such that (1) the row $\langle \#, r \rangle$ is not semantically determined by VAL_i, and (2) VAL_{i+1} is the proper subset of VAL_i that results when all the members in VAL_i that are not classically correct at $\langle \#, r \rangle$ are dropped from VAL_i. In other words, one obtains VAL_{i+1} from VAL_i by removing the semantic indeterminacy at r in $\#$. Finally, by the *length* of a semantic path we mean the number of non-initial members in the sequence, i.e., the number of members other than the initial element VAL_{snt}.

Semantic holism is the thesis that from inferential semantics to classical semantics there is but a single semantic step. That is, semantic holism maintains that the maximum *length* of semantic paths is 1. Clearly, if semantic holism were true, there would be just two semantic paths, one of length zero and the other of length one. The latter, of course, would end with VAL_{nrm} and so would generate classical semantics. More metaphorically, let us imagine a logician lost in the dark forest of indeterminate semantics. Then, according to the thesis of semantic holism, the logician need take but a single semantical step to escape from the dark forest of semantic indeterminacy.

Now that semantical holism has been shown to be false, one wonders whether semantical holism is *almost true* in the sense that there is a finite upper bound k on the length of semantic paths. In this case, all semantic paths would end at VAL_{nrm}, i.e., at classical semantics, if continued for at most k steps. So, to couch the question in the metaphor of the lost logician, is semantic holism almost true in the sense that the logician need take but a few steps to escape from the forest of semantic indeterminacy? Our first open question, then, is whether semantic holism is almost true.

Or, one might wonder whether semantical holism is *substantially true* in the sense that, although there is no finite upper bound on the length of semantic paths, there are nevertheless no infinite semantic paths. Then, if he or she would only take *enough steps*, our lost logician would be sure to escape from the semantical forest (i.e., to arrive at VAL_{nrm}). Our second open question, then, is whether semantic holism is substantially true?

Third, one might wonder whether there is at least *some truth* to semantic holism in the sense that, although there are infinite semantic paths, each of them converges on VAL_{nrm} (classical semantics) in the sense that, for each infinite semantic path P, the intersection of all the members of P is VAL_{nrm}. This would mean that our lost logician could count on escaping from the semantical forest (by arriving at VAL_{nrm}) after a literal eternity of wandering in the semantical forest. Our third open question, then, is whether there is some truth to semantic holism.

One wonders penultimately whether semantic holism is *seriously false* in the sense that there are infinite semantic paths that do not converge on VAL_{nrm}. This would mean that, even if he or she were to wander for a literal eternity, our lost logician might still fail to escape for the dark forest of semantical indeterminacy. This, then, is our penultimate question, whether semantic holism is seriously false.

We have run out philosophically interesting problems, but there will remain some fascinating technical questions even if semantic holism turns out to be seriously false. Their interest depends on the fact that there is only one intelligible way to extend the concept of semantic path into the transfinite: at the ordinal O, take VAL_{snt}; at a successor ordinal, take the result of removing a semantic indeterminacy (as above); and at a limit ordinal, take the intersection of all predecessors, which is the very set of valuations to which the predecessors converge. [In jargon terms, a *semantic path* is now any function P from an initial segment of the ordinals into sets of sound valuations such that (1) $P(0) = VAL_{snt}$, (2) if defined, $P(\alpha+1)$ is the result of removing a semantic indeterminacy relative to $P(\alpha)$, and (3) if α is a limit ordinal at which P is defined, $P(\alpha)$ is the intersection of all the sets $P(\beta)$ where $\beta < \alpha$.]

Therefore, even if it is seriously false, we can still sensibly ask if there remains some shred of truth at a transfinite distance from semantic holism. We are interested in possibly transfinite semantic paths that terminate in VAL_{nrm}, the determinate haven of classical semantics. We no longer need to worry about paths that converge on but do not terminate in VAL_{nrm}. Now that we have gone transfinite, we can just extend any such convergent path with VAL_{nrm} itself and thus create a path that strictly leaves the shadowy forest by terminating in VAL_{nrm}. Accordingly, we let a *normal semantic path* be any (finite or transfinite) semantic path of the kind just explained, provided that it has a least member and provided that this last member is VAL_{nrm}. We measure the *length* of such a path by the ordinal that marks its last member VAL_{nrm}, which is evidently convenient for us even if slightly nonstandard. So, a normal semantic path will eventually take the lost logician out of the murky forest, though the way may be long.

Let Π be the class of ordinals that mark the last member of some normal semantic path. Since normal semantic paths contain no repetitions, and since path membership is restricted to subsets of VAL_{snt}, evidently Π is not just some grossly bloated class, but rather is small enough to be a proper set of ordinals and so to have a least upper bound. This now becomes our focus. We let π be the ordinal that is uniquely determined as the least upper bound of Π. The problem of semantic holism is, then, this: How big is π, and (if it is a limit ordinal) is π a member of Π?

Why have we split the problem of semantic holism into two subproblems? First, we want to know the size of π because this will help us measure the degree of falsity of semantic holism, or if you prefer an equivalent image, the distance from semantic holism at which one finds a little truth. With regard to the second subproblem, notice that to say that π *is not* in Π is to say that, while every normal semantic path has length less than π, their lengths approach π without bound, whereas to say that π *is* in Π is to say that, while no normal semantic path is longer than π, nevertheless π is itself the length of some normal semantic path. The latter is clearly a more serious departure from semantic holism, in the same way that our concept above of "some truth" was further from semantic holism that "substantial truth". For, if π is not in Π, then every normal semantic path will take you out of the forest without your having to travel further than some ordinal strictly less than π; but if π is itself in Π, then at least one normal semantic path will force you to travel the full length π in order to get out. (If π is a successor ordinal, it is bound to be itself in Π, so that the second subproblem arises only when π is a limit ordinal.)

Our earlier list of open questions can be translated into the Π-and-π jargon thus. If semantic holism had been true, π would have been 1. If semantic holism is almost true, then π is finite. If semantic holism is substantially true, then π is ω, and is not itself a member of Π. And if semantic holism contains some truth, then π is ω, and is itself a member of Π.

Suppose, however, that semantic holism is seriously false in the sense of our penultimate question above, i.e., suppose that $\pi > \omega$. We can still pose the following questions. Is π countable or not? Is π a limit ordinal or not? Suppose π is a limit ordinal. Does π belong to Π or not?

What is the worst case for semantic holism? Begin by observing that there are only countably many connectives and therefore only countably many tabular rows. It follows that there can be only countable many successor stages in any normal semantic path, for each such stage must remove a semantic indeterminacy at some new row. However, since any uncountable ordinal is preceded by an uncountable number of successor ordinals, the last member of any normal semantic path must be marked with a countable ordinal. Therefore, the worst case for semantic holism would require that π be the first uncountable ordinal. Suppose that π is the first uncountable ordinal; what about the second subproblem? Even though π is a limit ordinal, we can rule out that it lies in Π, because, as we just saw, no normal semantic path can be marked with an uncountable ordinal. So, the *worst case* for semantic holism would occur if π should be the first uncountable ordinal, which (as just argued) does not lie in Π.

If the worst case does not obtain, so that there is some countable upper bound to the set Π of lengths of normal semantic paths, we might be willing to say that even if semantic holism fails to contain some truth in the sense of our third question, and thus is seriously false in the sense of our penultimate question above, still *its falsity is less than total.* But if it turns out that π is the first uncountable ordinal, so that there is no countable upper bound to the lengths of normal semantic paths — the worst case — we are certainly prepared to say that semantic holism is *utterly false.*

We close with a taxonomic challenge. We have divided tabular rows for connectives into those that are semantically determined and those that are not. Among the latter, we have seen that some rows represent semantic indeterminacies that are systemic, while others represent semantic indeterminacies that are not systemic. This suggests that there may be more ramified row-taxonomies that are as principled and at least as theoretically fruitful as our own simple classification of rows has proven to be. Is there anything to this suggestion?[6]

References

[1] R. CARNAP, ***Introduction to Semantics***, Studies in Semantics, vol. 1, Harvard University Press, 1942.

[2] R. CARNAP, ***Formalization of Logic***, Studies in Semantics, vol. 2, Harvard University Press, 1943.

[3] A. CHURCH, Review of ***Formalization of Logic*** by R. Carnap (Harvard University Press, 1943), ***The Philosophical Review***, vol. 53 (1944), pp. 493–498.

[4] G. MASSEY, *The pedagogy of logic: humanistic dimensions*, ***Teaching Philosophy***, vol. 4 (1981), pp. 303–336.

[5] J. D. MCCAWLEY, ***Everything that Linguists have Always Wanted to Know about Logic but were Ashamed to Ask***, The University of Chicago Press, 1981.

[6] D. J. SHOESMITH and T. J. SMILEY, ***Multiple-Conclusion Logic***, Cambridge University Press, 1978.

DEPARTMENT OF PHILOSOPHY
UNIVERSITY OF PITTSBURGH

Received July 12, 1988
Revised April 20, 1989

[6] We thank Stewart Shapiro for calling our attention to Carnap's 1943 book, which bears directly on what we call *the problem of semantic holism.*