Philosophical Papers

Volume I

DAVID LEWIS

New York Oxford
OXFORD UNIVERSITY PRESS
1983

Library of Congress Cataloging in Publication Data

Lewis, David, 1941–
Philosophical papers.

Includes index.
1. Philosophy—Addresses, essays, lectures.
I. Title.
B29.L49 1982 110 82-22551
ISBN 0-19-503203-9 (v. 1)
ISBN 0-19-503204-7 (pbk. : v. 1)

Printing (last digit): 9 8 7 6 5 4 3

Printed in the United States of America

· FOURTEEN ·

'Tensions

I have a problem for those who, like myself, admire intensional formal semantics and think it a key to understanding natural language. We can go to extremes of intensionality, if we like. Semantic rules can be stated entirely in terms of intensions, while extensions go unmentioned. But when we do, it seems for all the world as if we've gone purely *ex*tensional instead! Let me explain.

I. THE INTENSIONAL LANGUAGE L_I

I begin by describing an intensional language L_I. There are various categories of expressions.

S is the category of sentences. A sentence takes as its extension one of the two truth values, truth or falsehood. But the truth value of a sentence may depend on facts which vary from one possible world to another; on the time, place, speaker, and other features of context; on the values assigned to any free variables that may be present; on the resolutions of various sorts of vagueness; and perhaps on other things as well. So truth in L_I is a three-place relation. A sentence has a truth value as its extension relative to an *index,* that being a package of a world, a time,. . ., and whatever else (apart from meaning) might be needed to determine an extension. Equivalently, for each sentence we have a function from indices to truth values. For any possible index, the function gives the truth value of the sentence at that index. This extension-determining function is the intension (in L_I) of the sentence.

N is the category of names. A name takes as its extension the thing named—perhaps a concrete material object, perhaps some other sort of entity. Again, the extension may vary; so for each name we have a function from indices that gives the extension of the name at any possible index. This function is the intension (in L_I) of the name.

For any two categories X and Y, we have a third category X/Y of expressions which can combine with expressions of category Y to form compound expressions of category X. Examples: an S/N is something that can combine with a name to make a sentence, so it is an intransitive verb. An (S/N)/N can combine with a name to make an intransitive verb, so it is a transitive verb. An (S/S)/S can combine with a sentence to make something that in turn can combine with a sentence to make a sentence, so it is a dyadic connective or operator. There are infinitely many of these functor categories, though only finitely many are employed in L_I. An X/Y has no extension. Its intension is a function from appropriate intensions for members of category Y to appropriate intensions for members of category X. All the compositional semantic rules of L_I are given by one simple schema, with various categories put in for X and Y:

> If α is an X/Y with intension A and β is a Y with intension B, then the result of combining α with β is an X with intension A(B), the value of the function A for the argument B.

"Combining" may simply be concatenation; or it may be some more complicated operation, perhaps different for different functor categories or even for different members of one category. Or it may be that structures built up by many successive combinations are subsequently transformed as a whole to give the "surface" expressions of L_I.

I described L_I as an intensional language. More precisely, (1) in any case of compounding, the intension of the compound is a function of the intensions of the constituents, but (2) the extension of the compound is not always a function of the extensions of the constituents. Part (1) follows from the given form for semantic rules; part (2) is an additional stipulation.

This completes a partial description of L_I, in a style that should be familiar nowadays.[1] It will be useful to give a shamelessly idealized reconstruction of the way that style evolved.

[1] I have followed the treatment in my "General Semantics," in this volume, stripped of such frills as many-place functors and a basic category of common nouns. The use of Ajdukiewicz's functor categories makes for brevity in stating the rules. But I could make my point just as well by discussing one of Richard Montague's well-known treatments of natural language. See his "The Proper Treatment of Quantification in Ordinary Language," in Hintikka, Moravcsik, and Suppes, eds., *Approaches to Natural Language* (Dordrecht: Reidel, 1972), or Barbara Partee's paper in Milton K. Munitz and Peter K. Unger, eds., *Semantics and Philosophy* (New York: New York University Press, 1974).

In the olden days, we knew only of extensions. We used semantic rules of this form:

(1S) If α is an X with extension A, then the result of combining α with β is a Z with extension f(A).

with a special rule for each β; or, when generality was possible, of this form:

(1G) If α is an X with extension A and β is a Y with extension B, then the result of combining α with β is a Z with extension f(A,B).

Then we began to pay attention to languages in which extensions could depend on features of context. We relativized to various sorts of indices, but still adhered to the old dogma that the extension of the compound is a function of the extensions of the constituents. We had special and general rules of these forms:

(2S) If α is an X with extension A at index i, then the result of combining α with β is a Z with extension f(A,i) at index i.

(2G) If α is an X with extension A at index i and β is a Y with extension B at index i, then the result of combining α with β is a Z with extension f(A,B,i) at index i.

But we knew that in a few cases, the extension of a compound at one index could depend on the extensions of the constituents not just at that index but at other indices as well. That was so for variable-binding quantifiers if we took the assignments of values to variables as indices; and for modal and tense operators, taking the indices as worlds or times. So we learned to tolerate a few special nonextensional rules of the form:

(3S) If α is an X whose extension varies from index to index in manner A, then the result of combining α with β is a Z with extension f(A,i) at index i.

(For example: if α is a sentence which takes the truth value truth at all and only the worlds in a set A, then the result of prefixing $\Box$ to α is a sentence that takes the value truth or falsehood at world i according as A does or does not contain all worlds possible relative to i.) Often the manner of variation did not even need to be explicitly mentioned as an entity, still less stigmatized as an intension. In the case of quantification, at least, such rules were not considered a significant breach of extensionality.

But these special rules proliferated in number when we considered languages with many different modalities; and in variety when we admitted intensional predicate modifiers or intensional sentential operators that could not be handled as modalities with relations of relative possibility. Generality requires quantification over appropriate entities. Once these are at hand, it is natural to identify them with intensions, since they carry all needed information about the meanings of the corresponding expressions. Thus we progressed to rules of the form:

(3G) If α is an X with intension A and β is a Y with intension B, then the result of combining α with β is a Z with extension f(A,B,i) at index i.

When a rule specifies the extension at every index, it thereby specifies the extension-determining function that we have called an intension. It is only a short step, therefore, to rules of the form:

(4G) If α is an X with intension A and β is a Y with intension B, then the result of combining α with β is a Z with intension f(A,B).

But this short step gives us a new freedom. For intension-specifying rules, unlike any sort of extension-specifying rules, can apply even when the resulting compound belongs to a category for which there are no appropriate extensions—for instance, when it is a compound modifier or quantifier or connective. In such cases, of course, the appropriate sort of intension can no longer be an extension-determining function from indices. We have seen already, in connection with the functor categories of L_I, what else it might be.

When an intensional rule is needed, an extensional rule will not work. But when an extensional rule will work, an intensional rule also will work. If we need some intensional rules, then we may gain uniformity by using intensional rules throughout, even where extensional rules would have sufficed. (It is a matter of taste whether this gain outweighs the waste of using needlessly intensional rules. For me it does.) At this point, our compositional semantic rules always specify how intensions of compounds depend on the intensions of their constituents. The extensions have faded away. We have come full circle, in a way: once again expressions are assigned semantic values on a single level, not two different levels. My description of L_I exemplifies this final, purely intensional, stage in the evolution of formal semantics.

II. THE EXTENSIONAL TRANSFORM L_E

Now I shall describe another language L_E; this time, an extensional language specified by purely extensional rules of the forms (1G) and (2S) above. L_E and L_I are obviously not identical, since one is an extensional language and the other is not. But they are closely related. All the structure of L_I is mirrored in L_E. Following Terence Parsons (more or less), I shall call L_E an *extensional transform* of L_I.

L_E has a category of names; but we do best to divide this into two subcategories according to the sort of thing that is named. S is the first category of names; an S-name takes as its extension—that is, it names—a function from indices to truth values. After all, any old entity is entitled to bear a name; nameability is not a special privilege of concrete particulars! N is the second subcategory of names; an N-name takes as its extension a function from indices to entities of any sort.

(Although S-names are not N-names, an appropriate extension for an S-name would also be an appropriate extension for an N-name). Names, in either subcategory, have their extensions rigidly. Their extensions do not vary from one index to another. We could say that the intension of an S-name or an N-name in L_E is the function which gives, for each index, the extension of the name at that index. (These are functions from indices whose values are themselves functions from indices.) But since these intensions are constant functions, they are scarcely worth mentioning.

L_E also has infinitely many functor categories: for any suitable categories X and Y, there is a third category X/Y of expressions which can combine with expressions of category Y to form compound expressions of category X. The "suitable categories" are the subcategories S and N of names, and also the functor categories themselves; but *not* the two further categories of L_E that we have not yet mentioned. The extension of an X/Y is a function from Y-extensions to X-extensions. All but one of the compositional semantic rules of L_E are given by this general, purely extensional schema:

> If α is an X/Y with extension A and β is a Y with extension B, then the result of combining α with β is an X with extension A(B), the value of the function A for the argument B.

Functors also have their extensions rigidly in L_E. We could take their intensions in L_E to be the constant functions from indices to their unvarying extensions. But why bother?

Such are the lexica of noncompound expressions in L_I and L_E, and such are the combining operations associated with functors (alternatively, the transformational apparatus) in L_I and L_E, that L_I is related as follows to the fragment of L_E presented so far. Let us say that the categories S of L_I and S of L_E are *namesakes*; also N of L_I and N of L_E; also X/Y of L_I and X/Y of L_E, whenever the X's are namesakes and the Y's are namesakes. Then for every pair of namesake categories of the two languages, exactly the same expressions belong to both. Further, whenever an expression α has A as its *in*tension in L_I, then also α has A as it *ex*tensions in L_E.

So far, L_E scarcely deserves to be called a language, for a language needs sentences. It would be quite wrong to think that S of L_E is the category of sentences. Tradition makes clear that the extensions of sentences are supposed to be truth values; whereas the extensions of members of S in L_E, we recall, are functions from indices to truth values. These are appropriate *in*tensions, but inappropriate *ex*tensions, for sentences. They are appropriate extensions only for names.

L_E does have a category of genuine sentences, however; and also a category of predicates for use in forming sentences. But there is only one predicate: $\boldsymbol{\Sigma}$. That is our metalinguistic name for it; actually, it is written as a blank space and pronounced as a pause. The remaining compositional semantic rule of L_E, again a purely extensional rule but special to the predicate $\boldsymbol{\Sigma}$, is as follows:

If α is an S with extension A, then the result of prefixing $\mathbf{X}$ to α is a sentence having as its extension at an index i the truth value A(i) given by the function A for the argument i.

So $\mathbf{X}$ is a sort of truth predicate; but a safe one, since there is no way in L_E to produce a paradoxical diagonalization.

We may note that whenever α is a sentence of L_I, then $\mathbf{X}\ \alpha$ is a sentence of L_E with exactly the same truth conditions. There are no other sentences of L_E besides these images of the sentences of L_I.

This completes my description of L_E. It is an extensional language, as I said it would be. Inspection of the semantic rules confirms that the extension of a compound is always a function of the extensions of the constituents.

III. THE PARSONS TRANSFORM L_P

If we like, we can take the names and functors of L_E and trade them in for predicates and quantified variables. For instance, if we have

$$\mathbf{X}\,((\alpha(\beta))(\gamma))$$

as a sentence of L_E, where β and γ are in the category N and α is in the category (S/N)/N, we can replace it by

$$\exists v \exists w \exists x \exists y \exists z\ (\mathbf{X}\,v\ \&\ Rvzw\ \&\ Rwyx\ \&\ \bar{\alpha}x\ \&\ \bar{\beta}y\ \&\ \bar{\gamma}z)$$

where R is a predicate meaning "—is the result of operating on the argument--by the function . . . " and $\bar{\alpha}$, $\bar{\beta}$, $\bar{\gamma}$ are monadic predicates uniquely satisfied by the entities that are the extensions in L_E of α, β, and γ, respectively. In this way we go from L_E to another extensional language L_P (if you consider quantifiers extensional). I shall call L_P the *Parsons transform* of our original language L_I, since when Terence Parsons speaks of an "extensional transform" he means a language like L_P.[2] But L_E is just as extensional as L_P. There is nothing anti-extensional about names and functors per se. Eliminating them is one enterprise, going extensional is another.

IV. THE PROBLEM

L_I can be a richly intensional language, whereas L_E is strictly extensional. An important difference, as we all were taught. But almost the only difference there is between the two!

[2]See Parsons, *A Semantics for English* (duplicated, 1968), and "Some Problems Concerning the Logic of Predicate Modifiers," *Synthese* 21 (1970): 327–29.

Two field linguists, I and E, fully equipped to perpetrate Cartesian deviltry, go to work on a certain tribe. They investigate the dispositions to verbal behavior under a wide range of deceptive stimulation, the beliefs and desires that would rationalize that behavior, and the neural hookup and laws that would explain it materialistically. They study these things until there is nothing more to know. Then I announce my conclusion: these tribesmen use the intensional language L_I. My colleague E, a keen extensionalist, disagrees. He thinks it gratuitous of me to ascribe to them a language that requires the notoriously obscure apparatus of intensional semantics. After all, a better explanation lies close at hand! His opinion is that they use the extensional language L_E. Dumbfounded by E's perversity, I know not what to say.

V. BAD REJOINDERS

I really don't know. I do know of several unsatisfactory arguments against E's opinion, and we had better clear those out of the way.

First, I might try arguing that E's account is worse than mine just because it is more complicated. He requires two more categories, one more lexical item, and one more rule. Besides, his extra rule departs from the standard form of his other rules.

But this argument is bad for two reasons. For one thing, extensionality itself is generally thought to be an important dimension of simplicity. E may say that it is cheap at the price. For another thing, I agree with E that a complete account should mention that speakers pause (and writers leave extra spaces) at the beginnings of their sentences. E has already covered this fact in his ascribed syntax and semantics. I have not, and I must find a place for it, at some cost in complexity, elsewhere in my total description of the tribe's use of language.

Second, I might try arguing that E's opinion goes against our paradigm cases of extension-bearing. "Boston" names Boston, for instance, and does not rather name some function from indices.

The paradigms, however, are cases of extension-bearing in certain particular languages: German, Polish, English, and some other familiar languages that can be translated into these by well-established procedures. We have no paradigm cases of extension-bearing in the language of these hitherto unstudied tribesmen.

Even if, in my opinion, their language does happen to be one of the familiar ones, still E cannot be expected to agree that the paradigms apply. For E and I disagree about which language is theirs.

Tarski's Convention T and its relatives will not help. Since the tribe's language is not—not uncontroversially, anyway—the same as our metalanguage for it, the only versions of these principles that apply are the ones stated in terms of translation. For instance, E and I may agree that a metalinguistic sentence of the form "——— names ——— in their language" (or "——— is a name having

——— as extension in their language") should be true whenever the first blank is filled with a name (in our language) of some name α in the tribe's language and the second blank is filled with a translation of α into our language. This gets us nowhere. Disagreeing as we do about what the names are and what their extensions are, E and I have no business agreeing about what the correct translations are.

Third, I might try arguing that their language cannot be an extensional language, as E claims, because certain inference patterns are invalid in it that are valid in any extensional language. For instance, I might point to inferences by Leibniz's Law, or by Existential Generalization, in which true premises yield false conclusions.

E should agree with me that Leibniz's Law (for example) preserves truth in any extensional language. He should also agree with me that truth is not preserved in the inferences I produce as counterexamples.[3] But he should not agree with me that those inferences are instances of Leibniz's Law. An inference by Leibniz's Law needs an identity premise, and how do we identify those? Not by looking for a stack of two or three or four horizontal lines! Semantically, an expression with two gaps expresses identity iff (1) the result of inserting names in the gaps is a sentence, and (2) the sentence so formed is true if the inserted names are coextensive, otherwise false. An identity premise is a sentence formed by thus inserting names in the gaps of an expression that expresses identity. Since E and I disagree about which are the coextensive names, we will disagree also about which are the expressions that express identity, which sentences are identity premises, and which inferences are genuine instances of Leibniz's Law. If E identifies instances of Leibniz's Law correctly according to his opinions about names and their extensions, the inferences he selects will indeed preserve truth.

Fourth, I might try arguing ad hominem that E has not really managed to escape intensionality, since the things he takes for extensions are intensional entities. Functions from indices to truth values are commonly identified with propositions (especially if the indices consist of possible worlds and little else). Functions from indices to things in general are likewise identified with individual concepts. How can *in*tensional entities be *ex*tensions?

But this is confusion. Intensionhood is relational. Intensions are things that play a certain characteristic role in semantics, not things of a special sort. E and I agree that in a suitable language (whether or not it happens in the language of this tribe) the very same thing that is the intension of one expression is also the extension of another. For instance, speaking in a fragment of technical English suited for use as the metametalanguage of a smaller fragment of English, we agree that one and the same thing is both the intension of the object-language expression "my hat" and

[3]I am here ignoring our disagreement about whether an S must be preceded by Σ to make a sentence. Strictly speaking, if $\alpha, \beta, / \therefore \gamma$ is a non-truth-preserving inference in L_I, then $\Sigma\alpha, \Sigma\beta, / \therefore \Sigma\gamma$ is a non-truth-preserving inference in L_E. The original Σ-less version is not any sort of inference in L_E, since its "premises" and "conclusion" are S-names rather than sentences.

the extension of the metalanguage expression "the intension of 'my hat'." In itself, this thing is neither an intension nor an extension.

What is true is that some things can serve only as extensions, while other things—functions from indices, for instance—can serve either as extensions or as intensions. But there is no kind of thing that is ineligible by its nature to be an extension.

Fifth, I might try arguing that E's opinion ascribes an extravagant ontology to the tribesmen. When I say that a certain word of their language names a certain concrete, material hill, E says that it names something rather more esoteric: a set-theoretic object built up from a domain of individuals that contains unactualized possibilia.

E and I, if we are consistent, believe in these esoteric entities ourselves. We do not doubt that we can have names for them. Then on what grounds can I deny E's claim that the tribesmen also have names for them? In fact, we both agree that they have names for other entities far more suspect, to wit certain far-fetched gods (according to me) or functions from indices to such gods (according to E).

I might do better to argue that the ontology ascribed by E is bad not because certain esoteric things are present in it but rather because certain unesoteric things are missing from it. Saul Kripke has suggested (in conversation, 1972) that it is wrong to ascribe to someone an ontology that contains sets without their members, functions without their arguments and values, or the like.

A plausible principle. But has E really violated it in ascribing the use of L_E, a language in which all the names are names of functions from indices and none are names of the concrete, commonplace things that are among the values of those functions? I think not. The ascribed ontology is not the same thing as the ascribed set of name-bearers. If there is an ontology associated with our language, for instance, it includes all the real numbers; not just the countable minority of them that bear names. It is of no significance that the set of name-bearers violates Kripke's closure principle, unless it can be shown to be the whole of the ascribed ontology. But it is hard to say what ontology, if any, is ascribed in ascribing the use of L_E. One looks for the domain of quantification. But L_E has no quantifiers! Quantifiers are sentence-makers; but the only sentence-maker in L_E is Σ , and that is no quantifier. So L_E does not have a domain of quantification in any straightforward sense, either one that satisfies the closure principle or one that does not.

Unlike L_E, the Parsons transform L_P does have a natural domain. More precisely, there is a set D such that we get the intended truth conditions for those sentences of L_P that are transforms of sentences of L_I iff D is included in the range of the quantified variables. (This assumes that the predicates of L_P have their intended interpretations.) The set D, which is the same as the set of extensions of expressions in L_E, does violate Kripke's closure principle and so is unsuitable to be ascribed as someone's ontology. If some extensionalist claimed that our tribesmen used L_P, disguised by transformations, I think we would have a promising line of attack against him. But how does that affect E, whose claim is different? Perhaps there is some way to show that if it is bad to ascribe the use of L_P, then it is just as bad to ascribe the use of L_E. But so far, this looks to me like nothing better than guilt by association.

VI. COMMON GROUND

So I have no way to argue against E's absurd opinion. But though we certainly disagree about something, the extent of our differences should not be exaggerated. In some sense, we have given equivalent descriptions of the phenomena. (That is just what makes it so hard to build a case for one description against the other.) More precisely, if we treat our semantic jargon as theoretical vocabulary and eliminate it by Ramsey's method of existential quantification, then our disagreement will vanish. Our two accounts are equivalent in the sense that they have a common Ramsification.

We can agree on the following description in neutral terms. There is a system of categories: call them just S, N, and X/Y whenever X and Y are categories in the system. There are three relations of expressions to things: call them the 1-'tension, 2-'tension, and 3-'tension relations. The 1-'tension of an expression, but not the 2-'tension or 3-'tension, may vary from one index to another. The 1-'tension of an S at an index is a truth value; the 1-'tension of an N at an index may be anything; an expression in one of the other categories has no 1-'tension. The 2-'tension of an S or an N is the function from indices that gives the proper 1-'tension at each index. The 2-'tension of an X/Y is a function from appropriate 2-'tensions for members of Y to appropriate 2-'tensions for members of X; if α is an X/Y with 2-'tension A and β is a Y with 2-'tension B, then the result of combining α with β is an X with 2-'tension A(B). The 3-'tension of any expression is the constant function from indices to the unvarying 2-'tension. Finally, a tribesman speaks the truth in his language just when he utters an S preceded by a pause, and the 1-'tension of that S at the index (or set of indices) determined by the occasion of utterance is truth.

So far, so good. To complete my account, I need only add a gloss: the 1-'tensions are the extensions, the 2-'tensions are the intensions, the 3-'tensions are neither, and S is the category of sentences. To complete his contrary account, E need only add his contrary gloss: the 2-'tensions are the extensions, the 3-'tensions are the intensions, the 1-'tensions are neither, and S is a subcategory of names.

But these disputed additions do not add much. No matter which way we apply our traditional semantic vocabulary of extension, intension, naming, and sentencehood, the facts of the matter are already covered by the unglossed neutral description which is the common Ramsification of both our opinions. That is all we would need, for instance, in giving an account of the tribe's use of language as a rational activity for imparting information, or as a physical phenomenon. The questions under dispute are, so far as I can see, idle. If they are not, their import should give us a way to settle them.

VII. MORALS

This story has an abundance of morals. But most of them are available only if you recklessly conclude that because I have not been able to solve my problem, therefore it must be insoluble. Actually I believe nothing of the sort. Probably there is some perfectly good reason why L_I and not L_E is the tribe's language, and I have just overlooked it. Still, what if the problem were insoluble?

First moral: we would have another good example of Quine's inscrutability of reference. Different assignments of extensions would account equally well for the solid facts of the matter, and there would be nothing to choose between them. But at the same time the indeterminacy would be made to seem less formidable than we might have thought. In view of the common Ramsifications, it seems that even the semantic facts are not in dispute. E and I disagree about the proper way to describe those facts in our traditional semantic jargon. That is only a localized indeterminacy in one small region of our language. Can Quine's other alarming indeterminacies be disarmed in the same way?

Second moral: intension and extension would be correlative concepts. Neither would make sense except by contrast with the other. In a two-level semantic analysis, where expressions are assigned certain entities that are functions and other entities that are the values of those functions at particular indices, clearly the functions are intensions vis-à-vis the values and the values are extensions vis-à-vis the functions. But in a one-level analysis, whether we approach it by getting more and more intensional or less and less, there is no more contrast and the correlative terms might therfore be out of place.

Our *final moral* is unconditional. It stands whether or not my problem can be solved. So elusive is the difference between using an intensional language and using an extensional language that it can scarcely matter much which we do use. Those who value the superior clarity of extensional languages as such[4] are misguided. There are differences that do matter. There is the difference between languages that can be analyzed by the methods of formal semantics and ones that (so far as we know) cannot. For any sort of ontic purist, there is the difference between languages that can be analyzed without recourse to suspect entities and ones that cannot. And there is the difference between standard first-order predicate calculus and all less familiar and less well-investigated languages. But none of these real differences between better and worse languages coincides with the difference between extensional and intensional.[5]

[4]For instance, David Lewis when he wrote the opening lines of "Counterpart Theory and Quantified Modal Logic," in this volume, as follows: "We can conduct formalized discourse about most topics perfectly well by means of our all-purpose extensional logic. . . . Then we introduce modal operators to create a special-purpose, nonextensional logic. Why this departure from our custom?" He proceeded to put the departure right; but his views might just as well have been presented as an extensional semantic analysis of an intensional language.

[5]I am grateful to Graham Nerlich and Max Cresswell for the conversations in which this paper had its origins, and to the National Science Foundation for research support.