Nuel Belnap

Under Carnap's Lamp: Flat Pre-semantics

Abstract. "Flat pre-semantics" lets each parameter of truth (etc.) be considered separately and equally, and without worrying about grammatical complications. This allows one to become a little clearer on a variety of philosophical-logical points, such as the usefulness of Carnapian tolerance and the deep relativity of truth. A more definite result of thinking in terms of flat pre-semantics lies in the articulation of some instructive ways of categorizing operations on meanings in purely logical terms in relation to various parameters of truth (etc.); namely, closing vs. leaving open, local vs. translocal, and anchored vs. unanchored. Basic relations among these categories are established.

Keywords: semantics, pre-semantics, truth, meanings, operations, Carnap

1. Introduction

Semantics presupposes grammar.[1] There are nevertheless "pure theories" of values (such as the extensions of [3]) and meanings (such as Carnap's intensions) that are unencumbered by grammar, and that are in this sense properly *pre-semantic* rather than "semantic" in the strict sense. The only application of pre-semantics is to semantics itself, and all its conceptions are directed to this end. Why then isolate pre-semantic concepts? In the first place, pre-semantics helps us become clear that some of the deepest semantic ideas are quite independent of notational systems (grammars). Second, in the tolerant spirit of Carnap, we believe that one is likely to want a *variety* of complementary (noncompeting) pre-semantic analyses—and most especially, a variety of pre-semantic treatments of one and the same "language." One does not have to "believe in alternative logics" to repudiate the sort of absolutism that comes not from logic itself, but from narrow-gauge metaphysics or epistemology. Carnap tried to soften this absolutism by illustrating with his two "methods," and his variable language "L", but although Carnap's beneficent influence is legendary, it seems worth repeating the lesson: There can and should be multiple useful, productive, insightful and pertinent analyses of the *same* target. Pre-semantics therefore emphasizes the usefulness of thinking in terms of a *variety* of pre-semantic systems.

One thing that stands out more clearly in pre-semantics is the likeness between the semantic treatment of different semantic parameters as they arise

[1] Cheerful thanks to Martin Allen, Adrian Staub, and Matthew Weiner.

Presented by **Heinrich Wansing**; *Received* April 6, 2004

Studia Logica (2005) 80: 1–28
DOI: 10.1007/s11225-005-6774-6

in various branches of logic. A pre-semantic policy of help in this endeavor is reliance on "*flat*" pre-semantics, which is the topic of this essay. By this we mean a style of pre-semantics that lets each and every parameter of truth stand on its own, democratically, so that the individual contributions of each parameter—be it domain, interpretation of a particular predicate letter, assignment to a particular variable, set of worlds, etc.—can be fairly discussed without metaphysical or epistemic distraction. Indeed, our chief aim in using flat pre-semantics is to describe purely logical ways of categorizing the semantic meanings of "modes of combination" (e.g. connectives, operators on terms, predication) in their various relations to individual parameters in a fashion that ignores the unsteady boundary between "extensional" and "intensional" logic. See section 5 especially for the parameter-relative ideas of local vs. translocal, closing vs. leaving open, and anchored vs. unanchored, as they apply to the pre-semantic correlates of connectives and the like.

The flat pre-semantic approach makes it obvious that truth is nearly always relative. In contrast, sometimes philosophers speak in a way that presupposes that "the" fundamental notion of truth is absolute. If, however, "fundamental notion" means "the notion to which you should primarily pay attention," then that is the wrong track. In fact: The truth is seldom absolute. In helpful Tarski-style semantic analyses, *the fundamental concept of truth is hardly ever the (of course definable) absolute version.* It is instead almost always the version of truth that is relativized to (or parameterized by, or made to depend upon) something. Analogy: The parenthood binary relation is more fundamental than the (definable) one-place parenthood property, even if for a particular stated purpose the property is more "important" than the relation. In old-fashioned language, "parenthood" is an essentially relational concept; and so is truth.

The jargon of Tarski's own articles (such as [6]) tends to conceal this fact. Tarski's fundamental notion is "satisfaction," in the expression "sequence s satisfies sentence A." That, however, is but a stylistic variant of "sentence A is true with respect to the sequence s." The latter terminology emphasizes that Tarski's essential idea was to relativize truth—to sequences—whereas his own terminology, we think in some part for the worse, suppresses this fact. Flat pre-semantics allows us to throw a raking light on Tarski's insight, and thereby reveals what is all too easily ignored.

There is in Tarskian model theory additional relativization of truth: Even simple truth-table analysis relativizes truth, namely, to "rows" or "assignments" (the terminology is irrelevant). For example, you are not entitled to the concept of a "tautology" unless you are willing to speak words such as "A is 'true in every row'," where an English preposition, here "in," carries the relativization. It is always easy and sometimes useful to suppress this fact of practice; here, however, we let the fact shine forth in order to provide what illumination it can. Our special concern is with the massive likeness of the practices of "extensional" and "intensional" semantics.[2] It will be obvious, however, that we are proceeding at a level of generality that encourages us to see many important likenesses among semantic analyses, including a special favorite of Carnap, the likeness between the semantic treatment of sharply different "categorematic" grammatical categories such as sentences on the one hand and terms on the other, and between different "operator" grammatical categories such as sentential connectives, term operators, and predicates.

The chief technical results of this essay occur only in section 5; earlier sections lay the groundwork so that those results may be put in the proper context. Before section 5, everything should look like the same old thing, flatly expressed. By exploiting the generality that flatness permits, we make some aspects of some standard ideas stand forth a little more clearly.[3] In section 5 we move beyond mere flattening: We categorize meaning-operations in terms of their structural relations to the parameters on which truth depends. Only at that point do we find some new ideas arising out of flat pre-semantics.

2. Intrusions

We keep grammar and set-class theory in the background. Here we say an intrusive word about each.

Grammatical intrusion. At the abstract level that is relevant to our concerns, we think of a grammar as involving the following. *Categorematic expressions*, such as sentences or terms, with the idea that a semantics will then give a "value" of some kind to each categorematic expression. *Syncategorematic expressions*, such as "∼" or "&" or "(", which play a role in some grammatical operation. *Grammatical operations, or modes of combination or functors*, each of which is a (grammatical) function taking categorematic expressions as input, and producing a categorematic expression as output.[4]

---

[2] We use the historical and practical word, "practices," because, following [7], we have neither seen nor been able to imagine a way in which to formulate a compelling theoretical difference.

[3] Some ... some ... standard: These are modest aspirations. They recognize that flattening can also make some things more difficult to see than they were before, and they certainly do not include the hope that all semantic ideas either ought to be or can be represented in flat pre-semantics.

[4] The grammatical phrase, "mode of combination," although striking a chord, is not

Example: the operation which, given two sentential inputs A_1 and A_2, produces an appropriate "conjunction" of those two sentences, perhaps having the appearance "$(A_1 \,\&\, A_2)$".

It is often reasonable, and Carnapian, to see the division between categorematic and syncategorematic expressions as more a matter for the linguist than for the speakers. Given a population of speakers in the process of communicating, there is no more reason to expect a uniquely determined "structural manual" than there is to expect a uniquely determined "translation manual."[5] As always, however, there is no inference from "not uniquely determined" to "not uniquely reasonable," much less to "unhelpful."

Set-class intrusion. We need to say something about sets, classes, and types. We do not need to say much, because the essential ideas of semantics and pre-semantics do not seem to us to depend much on foundational distinctions. Still, when we say "set" we have in mind something like Zermelo-Fraenkel set theory, and when we say "class" we mean to suggest a collection that might, by its sheer size, outrun that theory. Then when we come to collecting subcollections, there are two cases. (1) If we are collecting subsets of a given set, then we know by ZF that the collection itself is a set. (2) If, however, we collect subclasses of a given class, the new collection will need to be at a higher type. In exactly the same way, when we come to functions, there are two cases. (1) If the domain is specified here as a set, of course the function itself can be a set in the usual ZF way. (2) If, however, the domain is specified here as a class, then we should expect the function to be "up" at a higher type. In short, please interpret us as consistent. It can't hurt.

At various points, grammar, pre-semantics, and semantics deal with functions from entities of kind K. Each of these functions will have a definite n-arity ($0 \leqslant n$), but its particular n-arity will hardly matter. We minimize distracting detail later if we agree now to treat all of these n-ary functions from entities of kind K as technically one-place. We may do this by letting *lists* of length n serve as (single) arguments for the n-ary functions in which we are interested. We let $\varnothing$ be the empty list, and we feel free informally to "identify," when convenient, a singleton list with its single occupant. Given a technically one-place partial operation on lists, we feel free to call it n-ary ($0 \leqslant n$) when all of its arguments are of length n.

as accurate as Curry's word "functor." A functor is just: a grammatical function. From this perspective, the negation connective is not the symbol $\sim$; it is the very function (a functor) that maps each sentence A into $\sim A$.

[5] The phenomenon of linguistic underdetermination does not, that is, start with "meaning." It is already present with "structure." To suppose otherwise is to use the forces of metaphysical or epistemological predisposition in order to deform logical insight.

3. Pure theory of value and meaning

With Carnapian tolerance in mind, we shall think of S as standing for a particular "pre-semantic system." That will give us a modest way of uniting with a single label ideas that belong together.

3.1. Basic "ontology" of flat pre-semantics

A particular pre-semantic system, S, can be understood as constituted by six ideas. Here is a list.

SIX FUNDAMENTAL PRE-SEMANTIC IDEAS. Choose S. To understand S as a pre-semantic system (in the present sense), one needs to grasp the following: (1) *S-values*; (2) *S-points*; (3) *S-parameters*; (4) the *S aux-function* and *S-auxiliaries*; (5) *S-meanings*; and (6) *S-operations.*

The ideas of S-values, S-points, S-parameters, and the S aux-function will be primitive, whereas the ideas of S-meanings and S-operations will be defined; but the defined ideas are nevertheless equally essential to the notion of a pre-semantic system and therefore are equally entitled to be on our list of "fundamental pre-semantic ideas."

A "pre-semantic" system, S, becomes "semantic" by relating its (4) "S-meanings" to categorematic expressions of a language, and its (5) "S-operations" either to syncategorematic expressions used in making grammatical combinations, or to the modes of combination themselves (whichever one prefers). While emphasizing the *theoretical* difference between pre-semantics and semantics, there is, as we have said, no point to pre-semantic systems without potential semantic application, and we shall therefore largely suppress the difference in treating preliminary *examples.*

3.2. S_1: an example

It will help to start with a familiar example intended to introduce the six ideas as "explicanda" before we become rigorous. (Nevertheless, in order to put these six ideas into perspective, it may be helpful to consult the more abstract Fig. 1 below.) The example, S_1, which is just a pre-semantics for a standard quantificational logic, comes in two parts. First we describe the grammar of the language that is the target of S_1, because although theoretically independent, in practice, pre-semantics needs to be guided by grammar. And then we describe S_1 itself as a pre-semantic system that illustrates the six listed ideas.

EXAMPLE 3.1. (Simple quantificational language: its grammar). The target of S_1 is one of the standard quantification grammars, namely, a grammar with (only) predicate letters and individual variables as "atoms." Predication in the standard grammar combines a predicate letter with a suitable list of variables to form an "atomic sentence." S_1, however, declines to consider either the predicate letter or the variables as categorematic components of the predication; they serve S_1 as mere syntactic auxiliaries (syncategorematica). As a consequence, the grammar of the language *relevant for* S_1 may be given as follows.

Categorematic expressions. (Only) *sentences* that are (1) either open or closed, and (2) either atomic or complex.

Our grammatical terminology here is not the most common.[6] More often one reads "formula" where we say "sentence," and "sentence" where we say "sentence containing no free variable." The more common terminology is all right for many purposes, as long as one realizes that by design or not, it tends to conceal what we here emphasize, namely, the relativization of truth! In fact, grammatically speaking, there is nothing "nonsentential" about either "it is a horse" or "x is a horse," and we may expect semantically that each of these two sentences (or call them formulas as long as it does not confuse you) receives a truth value—properly parameterized, of course.

Modes of combination (that is, grammatical operations). (1) For each "atomic" sentence, a 0-ary grammatical operation that generates exactly that sentence (from the empty list of arguments).[7] (2) A few standard truth functional connectives, such as the negation connective, which transforms a sentence, A, into $\sim A$. (3) For each individual variable, x, at least one quantificational connective, such as the universal quantifier for x, which transforms a sentence, A, into $\forall x A$.

An n-ary connective properly speaking is any function from n-ary lists of sentences into sentences, with perhaps the added conditions that outputs uniquely determine inputs, and that there are no infinitely descending decompositional chains (no ambiguity). It is a regrettable accident of familiar *artificial* languages that it is easy when explaining logical grammar to confuse connectives with certain symbols or symbol-patterns (syncategorematica). Perhaps, however, *spreading* the confusion arises out a desire to enrich the logical message with some physicalistic subtext. Fortunately this particular confusion does not stand in the way of learning truth tables and the like. It is on the other hand an illogical and harmful historical aberration that it is seldom made explicit that *quantifiers are non-truth-functional connectives.*

[6] Passages in smaller print are explanatory—we trust usefully so. Still, there is some sense in which these passages are redundant for the purely technical development.

[7] This is of course a bit of a technical fiddle; we allow it not so much because it is enlightening, but because it makes it easier, later, to say more gracefully some things that would otherwise be awkward.

EXAMPLE 3.2. (S_1: a pre-semantics for a simple quantificational language). The pre-semantic system, S_1, is targeted at the quantificational language characterized in Example 3.1. The six fundamental ideas of pre-semantics, when specialized to S_1, we elucidate as follows.

1. S_1-values. The S_1-values are just the truth values, T and F; there are no other S_1 values. This follows up the decision to treat only the sentences (both open and closed) as categorematic: The idea of an "S-value" in this jargon is exactly what could be attached to a grammatically categorematic expression.

In each application, the semanticist is forced to make numerous decisions, some of which are "don't-cares" and some of which he or she might wish to defend on philosophical or empirical or other grounds. For us, in this enterprise, *all* such decisions are don't-cares in the (limited) sense that when we offer a semantic system based on certain of these decisions, we do not mean to suggest by that alone that other decisions would somehow be faulty. In treating standard quantification, for example, one might well choose to treat the variables as themselves categorematic, giving rise to another pre-semantic system, $S_{1'}$, distinct from S_1, intended to apply to essentially the same language. In $S_{1'}$, in addition to truth values, individuals would also be $S_{1'}$-values. We ought to count as foolish anyone who thinks that there is something *wrong* with $S_{1'}$; but provided tolerance reigns, there is also nothing wrong with S_1 choosing to place the emphasis on sentential values by denying, *in the context of its particular analysis*, a categorematic status to individual variables. Each of S_1 and $S_{1'}$ is enlightening in its own way.

2. S_1-points. S_1-points encode whatever information is needed in order to determine a definite S_1-value, that is, a definite truth value, for the considered simple quantificational grammar. Namely, as we learned from Tarski, a S_1-point encodes the following. (a) a domain, (b) for each predicate letter, an appropriate subset of (or relation on) the domain, and (c) for each individual variable, a member of the domain. One may think of these three sorts of items, to which we shall repeatedly return, either as "components" or as "coordinates" of S_1-points.

More customary are the longer phrases *valuation point* or *point of evaluation*, and indeed we always mean that the short word, "point," should carry the intent of these longer phrases. We could as well have said "S_1-index" instead of "S_1-point," as does [5].

3. S_1-parameters. The S_1-parameters articulate the S_1-points into their separate components (or features or aspects or whatever). In flat pre-semantics, we are after *the most refined articulation.* We wish to keep separate track of each feature of an S_1-point that contributes to the truth values of sentences. Namely, to review the items listed under our story about S_1-points, the following: (a) the domain parameter, (b) a parameter for each predicate letter, and (c) a parameter for each individual variable. These are the items of which S_1 needs to track separately as it examines their influence on S_1-meanings and S_1-operations as defined below.

The contrast is with a system that has only three parameters, (a) a domain parameter, (b) an interpretation (which would be internally complex), and (c) an assignment (also internally complex). The style of flat pre-semantics is *flat* precisely in the sense that the separate parameters are in no way organized into a hierarchy. Certainly various S_1-parameters (e.g., the domain parameter vs. the F_2 parameter vs. the x_1 parameter) have different conceptual roles, and certainly the various S_1-parameters naturally group together under the headings (a), (b), and (c). That is what makes it tempting (and often useful) to speak of just three parameters instead speaking of the flat list of all the infinitely many individual items. Here, however, our emphasis is explicitly structural, and at the level of structure, each parameter, whether for domain or for F_2 or for x_1, is just: a parameter. By this flattening we mean, for present purposes, to emphasize likeness over difference and articulation over clumping.

4. The S_1 aux-function and S_1-auxiliaries.[8] In order to appreciate the role of the S_1 aux-function, make a picture in your head of a rectangular array rather like the "reference columns" normally found over to the left of a truth table (see Fig. 1).

First, imagine that the individual S_1-parameters (or their names) are written across the top of the reference columns. In this position, S_1-*parameters can serve as column headings*: the domain parameter, the F_1 parameter, …, the x_1 parameter, … . Because we are doing "flat semantics," each parameter should be written separately. We shall use "p" to range over S_1-parameters.

Second, imagine the S_1-points (or their names) as written down the left. In this position, S_1-*points can serve as row headings.* In contrast to the S_1-*parameters*, however, you will have in mind no "natural" names for the S_1-*points.* So just make up a few: v_1, v_2, … . The letter "v" will remind us that these are "valuation points," relative to which we may find a truth value for each sentence.

So now you have S_1-parameters p as column headings, and S_1-points v as row headings, of a rectangular array. How is the array to be filled in? That is the job of the S_1 aux-function.

The S_1 *aux-function* tells you what lies at the intersection of row and column. It tells you, for example, what the domain is at the S_1-point v_1, or what the interpretation of F_3 is at the point v_5, or what the value of x_7 is at the point v_2.

[8] We put up with ugly jargon in order to emphasize the role of these items.

These "intersection entities" are *auxiliary* to the semantics: They need be neither S_1-values nor S_1-meanings. For lack of a better idea, we therefore write Aux_{S_1} for the S_1 aux-function. Aux_{S_1} is defined for every S_1-point v and every S_1-parameter p, so that the domain of definition of Aux_{S_1} is definable in terms of the other fundamental pre-semantic ideas.[9] Aux_{S_1} remains a primitive, however, because it is only constrained by rather than definable by those other ideas: The various values of Aux_{S_1} are radically diverse, and we simply label these values as S_1-*auxiliaries.* We write $Aux_{S_1}(v, p)$ for the value of Aux_{S_1} at the S_1-point, v, and the S_1-parameter, p. These S_1-auxiliaries fill in the rectangular array for which the S_1-parameters serve as column headings and the S_1-points as row headings. The S_1-auxiliaries include sets (the domains), more sets (interpretations of e.g. a one-place predicate letter, F_1), relations (e.g. an interpretation of a two-place predicate letter, F_5), and "individuals" (e.g., an assignment to x_2).

The aux-function for the particular system S_1 has the following properties. The lettering refers to our conception of S_1-points as described under (2) above.

(a) If p_1 is the domain parameter, $Aux_{S_1}(v, p_1)$ is always a nonempty set: "the domain of v." (b) If p_1 is the parameter for any n-ary predicate letter, F, then $Aux_{S_1}(v, p_1)$ is an n-ary relation on "the domain"; that is, on "the domain of v." In still longer words, for each S_1-point v, if p_1 is the parameter for a n-ary predicate letter, F, and p_2 is the domain parameter, then $Aux_{S_1}(v, p_1)$ is an n-ary relation on $Aux_{S_1}(v, p_2)$. Investigations not needing the present level of abstraction, or that are not harmed by concealing the fact that "the domain" is itself an S_1-parameter, often say "the interpretation (or value) of F in v," and use a much shorter notation such as "$v(F)$". (c) When p_1 is the parameter for the individual variable, x, then $Aux_{S_1}(v, p_1)$ is a member of "the domain" (as spelled out just above). A frequent jargon is something like "the value of x on v," or just "$v(x)$".

S_1 puts no further "necessary" conditions on its points, its parameters, and its aux-function. In other words, to a first approximation, given any nonempty set, and no matter how you choose appropriate values for the other parameters, there is always a S_1-point such that the S_1 aux-function gives that set to the domain parameter, and also gives each other parameter its chosen value.

[9] Some pre-semantic systems are perhaps best understood as offering $Aux_S(v, p)$ as an only partial function. As a mere technical convenience, and with no philosophical point in mind, we may simulate this partiality by giving $Aux_S(v, p)$ some dummy auxiliary value when it would otherwise be undefined.

How does this "first approximation" need to be refined? Well, it is possible to be cagey or disputatious or merely worried about just which nonempty sets are covered by the quantifier word "any" as it occurs in our approximation, "given any nonempty set." So far as we know, as long as S_1 covers a great many nonempty sets, we always find that S_1 provides significant logical enlightenment.

5. S_1**-meaning.** An S_1-meaning is *defined* as a function from the S_1-points into the S_1-values. S_1-meanings, like S_1-values (i.e., truth values) belong to sentences: The S_1-meaning of a sentence shows how its S_1-value (its truth value) varies as one varies the items listed just above under (a), (b), and (c) of our account of S-points.

The phrase "S_1-meaning" is neither pretentious nor unpretentious, but just right—for S_1. The system S_1 itself is both interesting and also relatively impoverished, and in exactly the same way, so is the idea of S_1-meaning. Philosophers who detest meaning will, in a sort of pseudo-scientific or atheistic spirit, reject the phrase "S_1-meaning" as meaningless; but that rejection is an aberration not to be encouraged. Other philosophers, those who think that all meaning must be deep, will, in a pseudo-humanist or worshipful spirit, reject the phrase "S_1-meaning" as heretical; but these philosophers are equally to be discouraged. Persons of sound judgment do not confuse meaning with religion.

Observe that this example makes it plain that "S-meaning" is an entirely abstract logical idea. It is not the same as the more specific and perhaps more metaphysical idea of "intension," when that is defined as (something equivalent to) a pattern of values (e.g. truth values or individuals) as one varies the world-of-evaluation parameter throughout the set-of-worlds parameter (see section 6).

Pre-semantics demands that each S-meaning have a comprehensible internal structure. It is this that separates S-meanings so sharply from S-values. Flat pre-semantics meets this demand by rendering S-meanings as functions—from (possibly structureless) S-points into (possibly structureless) S-values.

In [3], each sentence acquired both an extension and an intension. In very much the same way, and with explicit dependence on Carnap, both S_1-values (truth values) and S_1-meanings attach to sentences. There is here, however, a critical point that Carnap did not make sufficiently clear: Although given a sentence, A, it makes "absolute" sense to speak of its S_1-meaning, the same is far from true for the way S_1-values (truth values) relate to A. For A does not have *any* truth value "absolutely." Instead, of course, in S_1 "the truth value of A" is always relative to some S_1-point, which encodes domain, interpretation, and assignment.

It can therefore be either helpful or misleading to say that in S_1, sentences have both a truth value and an S_1-meaning. There will be no problem as long as one keeps firmly in mind both (1) that S_1 gives sentences truth values *relative* to an S_1-point, and (2) that S_1 gives sentences S_1-meanings "absolutely."

6. S_1**-operation.** A S_1-operation is *defined* as a function from S_1-meanings to S_1-meanings. Just as S_1-meanings belong (only) to sentences, so S_1-

operations belong (only) to connectives of the grammar being considered (Example 3.1). The S_1-operation associated with a connective shows how, in the spirit of Frege, the S_1-meaning of the constructed sentences depends (entirely) on the S_1-meanings of its sentential parts. In the grammar being considered, the primitive connectives include not only the truth functional ones, but also $\forall x_1$, $\exists x_2$, etc. It is obvious from Tarski's work that the S_1-meaning of e.g. $\forall x_1 A$ depends exclusively on the S_1-meaning of A, and that we understand the meaning of $\forall x_1$—if in fact we do—by understanding how it maps each possible S_1-meaning of the part, A, into a S_1-meaning of the complex, $\forall x_1 A$.

Spirit of Frege? Well, we do take up the Fregean idea that we think valuable, which is that "the meaning of a compositional whole should depend on the meanings of its parts." And we refuse to take up the Frege idea that "the truth value of a compositional whole should depend (entirely) on the truth values of its parts." The latter is violated in S_1 by the quantificational connectives, a point to which Fregean "senses" are irrelevant. One can of course present a semantic system for quantificational logic in a way that conceals that the meaning of e.g. $\forall x_1 A$ depends on the meaning of A, a procedure that is likely to be harmless where the aim is purely technical.

This finishes our account of the pre-semantic system, S_1. Before proceeding, and in order to encourage verbal explanations to mingle with geometric intuitions, we call attention to the diagram of Fig. 1, which is relevant to *any* pre-semantic system, S, whether it be S_1, or truth-tables, or a sophisticated intensional system. Speaking in terms of pictures, a flat pre-semantics *always* looks very much like a truth table.

3.3. Alternatives to the S aux-function

Because the idea of the S aux-function is clumsy, we mention two alternatives.

(1) One may identify S-points as functions from the S-parameters into the S-auxiliaries (or, if the parameters be few enough, one may identify them with sequences of S-auxiliaries *á la* Tarski). In this representation, which is perhaps most usual, only the S-parameters are taken as technically without internal structure, and one may drop the aux-function.

(2) One may identify each S-parameter as a function from S-points to S-auxiliaries. In this case, only the S-points are taken as technically without internal structure, and, again, one may drop the aux-function.

Plans (1) and (2) are felt to be substantially equivalent to each other; and each to be equivalent to a third plan, which is used here:

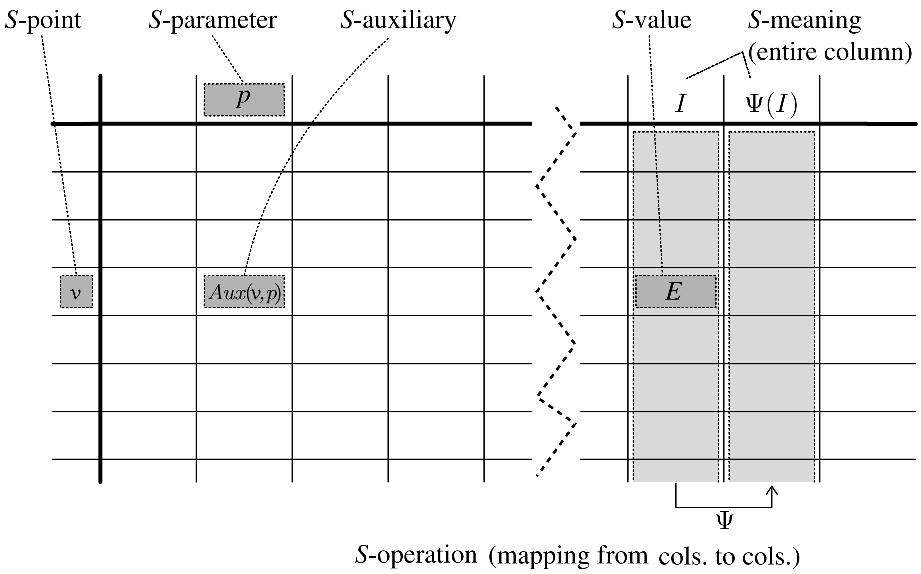

Figure 1. Six fundamental notions of pre-semantics

(3) S-points and S-parameters are each taken by S as without structure, leaving their interaction to be determined externally by Aux_S, the S aux-function. We use this less graceful third alternative not for novelty, but only to emphasize that connection between points and parameters does not depend on some set-theoretic trick. You must have S-points and you must have S-parameters, and whether or not points or parameters have a function-like structure, you must have an account of "the auxiliary determined by the S-point v and the S-parameter p." Flat semantics emphasizes this. For any other purpose, one should certainly feel free to use whichever of the three plans seems most helpful for the immediate purpose.

What does it cost us to be sure that the three plans are indeed equally available for S? Only that (1) S-point-as-function is really enough to identify a S-point, and (2) analogously for S-parameter-as-function. That is, we must say that S does not admit two (distinct) S-points whose structuring by the S-parameters is exactly the same, nor analogously for two (distinct) S-parameters. We'll refer to this property of a pre-semantic system as "*articulation.*"

The *benefit* is that if a pre-semantic system is articulated, we are less likely to be confused in our thinking about it just because of the equivalence of (1), (2), and (3). Evaluation of the *cost* of enforcing articulation depends on which direction we are considering. Half of articulation is that no two (distinct) parameters should behave exactly the same on all points. It is easy to invent a system that fails to have this feature. For example, let "gorse" and "furze" be two predicate parameters, and suppose that as a way of instituting a Carnapian "meaning postulate," we wished to constrain points such that the two predicate parameters must always be assigned exactly the same subset of the domain. It is hard to find any *logical* reason to forbid such a system even though it would not be articulated. For an example in the other direction, take the [4] case of contexts of utterance, and let these contexts be, for this toy example, S-points. Let the S-parameters be speaker, place, and time of context. Certainly speaker, place, and time do not exhaust all of the features of a context that are of interest to a philosopher of language. Consider that to process the indexical "you," one would wish at least to consider an audience or auditor of the context. It is therefore natural to imagine that S should admit two contexts as distinct points (or parts of points) that are exactly alike with respect to the given S-parameters, speaker, place, and time, but differing in e.g. auditor. In such a case, S would not be articulated, and you would be misled if you tried the easy maneuvering between (1), (2), and (3) that articulation underwrites.

3.4. Technical definition of a flat pre-semantic system, S

Just for the record, or, as Kaplan says, for a reality check, we put the key ideas together as a theoretical definition of a "pre-semantic system."

DEFINITION 3.3. Pre-semantic system S is a *pre-semantic system* iff S is a structure $\langle E_S, V_S, \mathcal{P}_S, I_S, \Psi_S, Aux_S\rangle$ satisfying the following conditions.

1. E_S is a nonempty class. Members of E_S are called *S-values.*

If we wish to deal with values appropriate to sentences classically understood, T and F, the two (distinct) truth values, presumably belong to E_S. If we wish to deal with values appropriate to singular terms, then very likely an extremely large array of individuals will figure as S-values. All that is required of an S-value is that it be thought of as a potential semantic value of a categorematic expression of a language at which S is targeted. There is no general theory of S-values as such.

2. V_S is a nonempty class. Members of V_S are called *S-points.* When S is understood, v ranges over V_S, so that v is a S-point.

An S-point is endowed with all the information required by S to determine a definite S-value for each categorematic expression in the target language. The system S_1 described above offers one familiar example of S-points. There is no general theory of S-points as such.

3. $\mathcal{P}_S$ is a nonempty class. Members of $\mathcal{P}_S$ are called *S-parameters.* When S is understood, p ranges over $\mathcal{P}_S$, so that p is a S-parameter; P ranges over subclasses of $\mathcal{P}_S$, so that P is a class of S-parameters; and $-P = \mathcal{P}_S - P$.

Technically speaking, a parameter is nothing but an index, a column heading of something required to fix a S-value. In each application, however, each parameter is given some significance that in practice is typically indicated by its name, e.g. "the domain parameter" or "the x_3 parameter," or "the set-of-worlds parameter." Indeed, because of the indexing role of parameters, as long as one knows what one is doing, there is no harm "identifying" a parameter with its name, or with the bit of notation, if any, with which it may be associated, or with a number or numeral picking out its position in some ordered list, if there be one. For example, it matters not if one identifies the x_3 parameter with "the x_3 parameter", or with x_3, or with "x_3", or with 3, or with "3." As long as its *function* is clear, its "*identity*" is of no consequence. There is no general theory of parameters as such.

4. For S a pre-semantic system, Aux_S is a (higher type) function whose domain of definition is exactly the class $V_S \times \mathcal{P}_S$. $Aux_S(v, p)$ is an *S-auxiliary value*, namely, *the S-auxiliary value determined by v and p.* When S is understood, p_v is defined as $Aux_S(v, p)$, and is read: the *value of* (the S-parameter) p at (the S-point) v. $Aux_S(v, p)$ is a *component of* (the S-point) v.

Auxiliary values need not (but may) also be S-values. For example, standard propositional logic takes sentences as its family of categorematic expressions, and awards truth values T and F to these as their S-values. In addition, however, a pre-semantics for this logic can use T and F as *auxiliary* values in the "reference columns," so that T and F serve a double role, as both S-values (of categorematic expressions) and S auxiliary values (of parameters). Here is a contrasting case. Standard model-theoretic quantificational semantics uses nonempty sets as auxiliary values of "the domain" parameter. These nonempty sets do not, however, occur as values of any categorematic expression. Auxiliary values need not (but may) attach to some "atomic" bit of the target language.

Articulation defined (but not assumed). For S-points v_1 and v_2, if $Aux_S(v_1, p) = Aux_S(v_2, p)$ for every S-parameter, p, then $v_1 = v_2$. For S-parameters p_1 and p_2, if $Aux_S(v, p_1) = Aux_S(v, p_2)$ for every S-point, v, then $p_1 = p_2$.

5. I_S is the (higher type) class $V_S \mapsto E_S$ of all functions from V_S into E_S. Members of I_S are called *S-meanings.* When S is understood, I ranges over I_S.

Later we describe a number of examples. The word "meaning" alone would here be Wrong, but "S-meaning" is honest and accurate, and there is, we think, no other word that will do. The intent of the pre-semantic system, S, is that S-meanings shall be the richest meaning that S attaches to categorematic expressions such as sentences or terms.

There is no assumption, however, that each S-meaning should attach to some categorematic expression; there are far too many S-meanings to expect that kind of expressive completeness.

6. 1-Ψ_S is the (yet higher type) class $I_S \mapsto I_S$ of all 1-ary functions from I_S into I_S, and similarly for n-Ψ_S. Members of n-Ψ_S are called *S-operations.* When $n = 1$, we write just Ψ_S. When S is understood, ψ ranges over Ψ_S.

Suppose $\psi \in \Psi_S$, $I \in I_S$, and $v \in V_S$. Then $\psi I \in I_S$, and $\psi I v \in E_S$. In designing a system, S, one intends to attach *pre-semantic* S-operations to *grammatical* modes of combination, those by which the target language constructs its complex categorematic expressions from its simpler categorematic expressions; for example, one may attach a certain S-operation to the negation connective.

3.5. Useful notation

We close this section by introducing some notation generally useful in discussing any pre-semantic system, S.

Definition 3.4. (Projection, agreement, parameter shift). Fix S. *Projection.* p_v (the projection of v onto p) is the value of the S-parameter, p, in the S-point, v. In other words, $p_v = Aux_S(v, p)$. *Agreement in P.* For any set, P, of S-parameters, $v_1 =_P v_2$ iff v_1 and v_2 agree on every S-parameter in P: $\forall p[p \in P \rightarrow p_{v_1} = p_{v_2}]$. Frequent case: $(v_1 =_{\mathcal{P}_S - P} v_2)$, which says that v_1 and v_2 agree everywhere *outside* of P. *Parameter shift.* Suppose that p is a S-parameter, v_1 is an S-point, and z any entity. Provided there is an S-point, v_2, such that (1) $v_1 =_{\mathcal{P}_S - \{p\}} v_2$ and (2) $p_{v_2} = z$, we define $[z/p]v_1$ as the S-point, v_2 such that (1) $v_1 =_{\mathcal{P}_S - \{p\}} v_2$ and (2) $p_{v_2} = z$. That is, $[z/p]v_1$ is the S-point that is exactly like v_1, except with the auxiliary value of the parameter, p, shifted to z (if there is such an S-point). When (1) and (2) hold, we say that $[z/p]v_1$ *exists.*

4. Properties of S-meanings in flat pre-semantics

Here we define, in the present setting, a standard way of categorizing S-meanings (here, S-meanings). Because the standard way is so thoroughly well-known, we may be extremely brief. In contrast, ways of categorizing S-operations have been little explored, which accounts for the greater length of the section that follows this.

Definition 4.1. (Properties of S-meanings). Fix S. I is *closed in* (or *constant in* or *categorical in* or *independent of* or *stable in*) P (or I is P-*closed*) iff $\forall v_1 \forall v_2[(v_1 =_{\mathcal{P}_S - P} v_2) \rightarrow (Iv_1 = Iv_2)]$. Otherwise I is *open in* (or *dependent on*) P. I is *(absolutely) closed* (etc.) iff I is closed (etc.) in $\mathcal{P}_S$. Otherwise I is *(absolutely) open* (etc.).

These various phrases, historically used in differing contexts and with differing rhetorical forces, have exactly the same structural meaning: namely, that the S-meaning, I, is such that, for each v, Iv will maintain a constant value no matter how you vary the P-components of v (as long as you leave the other components alone). In quantification theory, if x_1 does not (as a grammatical fact) occur free in A, then the S_1-meaning attached to A is (as a pre-semantic fact) certain to be closed in the parameter for x_1.

If I is absolutely closed, then I is of course a constant function, delivering always the same S-value at every S-point. The least interesting instance of this abstraction is the one most frequently exploited by logical theorists: When $Iv = \mathrm{T}$ for every v, then I is a pre-semantic representation of the S_1-version of "logical truth," a notion that many think is of too little logical utility in proportion to the attention that it has attracted.

5. Properties of S-operations in flat pre-semantics

Here we finally make good on our plan, adumbrated at the beginning of this essay, to use flat pre-semantics as a platform for describing purely logical ways of categorizing various ways in which S-operations can be related to S-parameters.

5.1. Operation properties explained

We categorize S-operations, in relation to a set of parameters, in terms of four fundamental dichotomies: essentially 0-ary vs. +-ary, local vs. translocal, closing vs. leaving open, and anchored vs. unanchored. These four dichotomies describe (not meanings but) S-operations in their relation to parameters. Doubtless these simple ideas have been isolated in similar semantic generality elsewhere, but, except for the first, we have not happened to come across them; therefore, unlike the dichotomous properties of S-meanings of Definition 4.1, we shall, regrettably, need to introduce unfamiliar words for them. We begin with a rough explanation of each of the four, specialized to the case of a single parameter (instead of a set of parameters). Also these preliminary explanations will concern the "semantics of connectives," a topic that is more familiar than the "pre-semantics of S-operations."

Four fundamental dichotomous relations of S-operations to S-parameters. The following thumb-rules, stated in terms of (relativized) truth, are intended to help explain the four dichotomies in a rough way. Imagine that we are considering an S-operation that is attached to some connective, Δ. Here, however, we loosen our account by speaking directly of the connective, Δ. The discussion pretends, in effect, that each S-meaning is expressed by some sentence, A.

We are trying to decide how Δ relates to some S-parameter, p.

Essentially 0-ary vs. essentially +-ary. Is ΔA always the same S-meaning regardless of A? If so, Δ is essentially 0-ary. But if sometimes A makes a difference to ΔA, then Δ is essentially +-ary.

Locality vs. Translocality. In calculating whether ΔA is true at an S-point, v_1, you will in general need to look at S-points other than v_1. But do you need to look at any that differ from v_1 on the S-parameter, p? If you never do, Δ is local in p. If sometimes you do, Δ is translocal in p.

Closing vs. Leaves open. Is the variation of p, taken by itself, irrelevant to the truth value of ΔA? In other words, if you hold all other parameters fixed and vary just p, does this ever make a difference to ΔA? In still other words, are there two S-points that are the same everywhere else but at p, and that nevertheless give different truth values to ΔA? If so, then Δ leaves p open. If not, Δ closes p. In other words, Δ closes p if ΔA is always closed in p, and Δ leaves p open if ΔA is sometimes open in p.

Anchored vs. Unanchored. Suppose that you can find a sentence, A, such that making a particular change in—and only in—p makes no difference to the pattern of values of A. But suppose that same change nevertheless makes a difference to the value of ΔA. If so, then Δ has the special "anchoring" relationship to p: In order to determine the value of ΔA, you (sometimes) need to know the very identity of the auxiliary value of p, and not just its contribution to the pattern of values of A. Otherwise, Δ is unanchored in p.

Here are the strict pre-semantic definitions of these four ideas.

Definition 5.1. (Essentially 0-ary vs. +-ary). Fix S. ψ is *essentially 0-ary* iff $\forall I_1 \forall I_2 \forall v[\psi I_1 v = \psi I_2 v]$; which is to say, iff $\forall I_1 \forall I_2[\psi I_1 = \psi I_2]$. In other words, and perhaps most usefully, $\exists I_2 \forall I_1[\psi I_1 = I_2]$. ψ is *essentially +-ary* iff ψ is not essentially 0-ary.

Unlike the dichotomies to come, the 0-ary vs. +-ary dichotomy does not relate to a specified set of parameters. For ψ to be essentially 0-ary is for it entirely to ignore its arguments; that is, an essentially 0-ary S-operation is a constant function, so that for any S-meaning argument whatsoever, its output is one and the same S-meaning.

Definition 5.2. (Closes vs. leaves open). Fix S. ψ *closes* P iff $\forall I \forall v_1 \forall v_2[v_1 =_{\mathcal{P}_S - P} v_2 \rightarrow \psi I v_1 = \psi I v_2]$. So ψ closes p iff $\forall I \forall v \forall z[[z/p]v$ exists $\rightarrow \psi I([z/p]v) = \psi I v]$. And ψ *leaves* P *open* iff ψ does not close P.

For an S-operation, ψ, to close a set of S-parameters, P, is for its every output, ψI, to be itself closed (or constant) in P.

A paradigm is the S_1-meaning, call it ψ, attached to $\exists x_1$. ψ closes the x_1 parameter. Take any sentence, A. Recall that in S_1, the truth value of a sentence depends on the domain, on the interpretation of each predicate letter, and on the value of each variable. That ψ closes the x_1 parameter implies that no matter the S_1-meaning, I, attached to A, the S_1-meaning, ψI, which is attached to $\exists x_1 A$, will be such that the truth value of $\exists x_1 A$ is certain *not* to depend on the x_1 parameter. If $\exists x_1 A$ is true [or false] at a certain point, v_1, then it will remain true [or false] if the x_1 parameter, which is a component of v_1, is varied in any way that you like. In other words, if $\psi I v_1 = \text{T}$ [or $= \text{F}$], then the same holds if v_1 is replaced by any other S_1-point that differs from v_1 only in respect of the value of its x_1 parameter.

An essentially 0-ary S-operation closes (leaves open) P iff the S-meaning that is its output is closed (open) in P.

DEFINITION 5.3. (Local vs. translocal). Fix S. ψ is *local in P at v_1* iff for every I_1 and I_2, if $\forall v_2[v_1 =_P v_2 \rightarrow I_1 v_2 = I_2 v_2]$ then $\psi I_1 v_1 = \psi I_2 v_1$. ψ is *translocal in P at v_1* iff there is a I_1 and a I_2 such that $\forall v_2[v_1 =_P v_2 \rightarrow I_1 v_2 = I_2 v_2]$ but nevertheless $\psi I_1 v_1 \neq \psi I_2 v_1$. That is, iff ψ is not local in P at v_1. We may also say that v_1 *witnesses the translocality of ψ in p*. ψ is *local in P* iff ψ is local in P at every v. ψ is *translocal in P* iff ψ is not local in P. So ψ is translocal in p iff $\exists I_1 \exists I_2 \exists z[\forall v[[z/p]v \text{ exists} \rightarrow I_1([z/p]v) = I_2([z/p]v)]$ and $\exists v[[z/p]v$ exists and $(\psi I_1)([z/p]v) \neq (\psi I_1)([z/p]v)]]$. Finally, ψ is *(absolutely) local [translocal]* iff ψ is local [translocal] in $\mathcal{P}_S$.

That is, ψ is local in P at v if it treats two S-meanings the same at v as long as those S-meanings are the same for S-points that agree with v inside of P. Thus, ψ is translocal in P at v iff you can find two S-meanings, I_1 and I_2, that are exactly alike for all S-points agreeing inside of P with v, but nevertheless $\psi I_1 v$ differs from $\psi I_2 v$.

Locality of ψ in P appears to have an intricate definition, but the idea is simple: In calculating the S-value of ψI at v, one needs to look at most at local auxiliary values of parameters in P. One does not need to look at any auxiliary P-values beyond the ones that occur as components of v. In contrast, if sometimes one needs to look at auxiliary P-values other than those occurring as components of v, then ψ is translocal.

Here are some examples from S_1 and nearby.

1. Let ψ be attached to any truth functional connective. Then ψ is absolutely local, hence local in *every* S_1-parameter.
2. Conversely, if any S_1-operation, ψ, is absolutely local, then ψ is truth-functional.
3. Switching grammatical categories, identity when added to standard quantification theory is absolutely local, just like truth functions.
4. Let ψ be attached to the connective, $\exists x_1$.
    - ψ is *translocal* in the x_1 parameter. This is an abstract way of noting that you cannot calculate a truth value for $\exists x_1 A$ at an S_1-point, v_1, without considering the various values of A at S-points, v_2, that *differ* from v_1 with respect to the x_1 parameter.

- But ψ is *local* in the x_2 parameter. That is, as Tarski taught us, whereas you must vary the auxiliary value of the x_1 parameter in order to calculate a truth value for $\exists x_1 A$ at v_1, at the same time you must not vary the value of the x_2 parameter. You must consider only other S_1-points, v_2, that are *exactly the same* as v_1 in their x_2 components.

5. One might have supposed that some combinations of closing vs. leaving open and local vs. translocal were ruled out; but in fact all are possible. In the following examples, permit us to omit some words by attributing these properties directly to the connectives themselves (instead of to the S_1-operations that the Tarskian semantics uniquely attaches to the connectives).

   - Negation is local in each S_1-parameter, and leaves each of them open. Ditto for $\forall x_1$ in the x_2 parameter.
   - $\forall x_1$ is translocal in the x_1 parameter, and closes it.
   - The connective that transforms A into $(\forall x_1 A \vee {\sim}A)$ is translocal in the x_1 parameter (because of the left disjunct) and also leaves open the x_1 parameter (because of the right disjunct).
   - The "excluded middle" connective taking A into $(A \vee {\sim}A)$ is local in every S_1-parameter, and also closes them all. (This connective is apparently unary, but essentially 0-ary.)

The purpose of these four examples is just to bring out that the terrain is too tricky to be hurried over. See Fact 5.6 and Fact 5.7 below for a more systematic survey. We remark in addition that if ψ is any one of the S_1-operations attached to a standard connective of first-order logic, then ψ is local in the domain parameter and in each predicate-letter parameter. And indeed it is precisely this fact that guides the formulation of the Tarski inductive definition of (relativized) truth, which holds the values of the domain and the predicate-letter parameters fixed, while allowing the (auxiliary) values of the individual-variable parameters to vary, by means of the distinctive Tarski concept of "satisfaction."

Finally, note that every essentially 0-ary S-operation is, vacuously, absolutely local.

Definition 5.4. (Anchored vs. unanchored). Fix S. ψ is *anchored* in P iff there is a S-meaning, I, and there are S-points v_1 and v_2, such that $v_1 =_{\mathcal{P}_S - P} v_2$ and $\forall v_3 \forall v_4[(v_3 =_{\mathcal{P}_S - P} v_4$ and $v_1 =_P v_3$ and $v_2 =_P v_4) \to Iv_3 = Iv_4]$ and $\psi I v_1 \neq \psi I v_2$.[10] So ψ is anchored in p iff $\exists I \exists z_1 \exists z_2 [\forall v[[z_1/p]v$ exists and $[z_2/p]v$ exists $\to I([z_1/p]v) = I([z_2/p]v)]$ and $\exists v[[z_1/p]v$ exists and $[z_2/p]v$ exists and $(\psi I)([z_1/p]v) \neq (\psi I)([z_2/p]v)]$. ψ is *unanchored* in P iff ψ is not anchored in P: For every S-meaning, I, and for all S-points, v_1 and v_2, if $v_1 =_{\mathcal{P}_S - P} v_2$ and $\forall v_3 \forall v_4[(v_3 =_{\mathcal{P}_S - P} v_4$ and $v_1 =_P v_3$ and $v_2 =_P v_4) \to Iv_3 = Iv_4]$, then $\psi I v_1 = \psi I v_2$. Or contrapositively: For every S-meaning, I, and for all S-points, v_1 and v_2, if $v_1 =_{\mathcal{P}_S - P} v_2$ and $\psi I v_1 \neq \psi I v_2$, then $\exists v_3 \exists v_4[(v_3 =_{\mathcal{P}_S - P} v_4$ and $v_1 =_P v_3$ and $v_2 =_P v_4)$ and $Iv_3 \neq Iv_4]$.

---

[10] The fundamental intuitions were worked out with Matthew Weiner, who provided the language of "anchoring," and with Martin Allen.

Again the definition is not immediately transparent, but the principle, although less familiar than e.g. locality, is important. Consider a S-operation, ψ, and a set of S-parameters, P. The question is, in calculating ψI, can you find a case in which a change in P does *not* make a difference to the pattern of values of I (as one varies parameters outside of P), but that same change *does* make a difference to ψI? If so, then ψ is anchored in P, and otherwise it is unanchored.

Example of anchoring. Consider the S_1-operation attached to $\forall x_1$, and consider the *domain* parameter. It "should" be that the universal quantifier on x_1 is anchored in the domain. And it is. It is sufficient to find a certain sentence, A, such that (1) its truth value is independent of the domain, whereas (2) the truth value of $\forall x_1 A$ depends on the domain. For example, let A be $F_1 x_1$. Evidently the truth value of this does not depend on the domain parameter (once every other parameter is fixed), whereas the truth value of $\forall x_1 F_1 x_1$ obviously does depend on the domain parameter (even if every other parameter is fixed). The input is independent of the domain, the output is dependent on the domain. So this single example is *more* than enough to show that the S_1-operation attached to the $\forall x_1$ is *anchored* in the domain parameter.

Example of unanchoring. Consider again the S_1 operation attached to $\forall x_1$, but now consider the F_1 parameter. It "should" be that the universal quantifier on x_1 is unanchored in the F_1 parameter. The idea is that in passing from A to $\forall x_1 A$, you cannot find a pair of S_1-points such that a difference in truth values of the output at those points can be attributed entirely to a change in the (auxiliary) value of F_1, except insofar as that change contributes to a change in the pattern of values of the input, A.

Perhaps it will help if we reduce to absurdity the claim that the universal quantifier in x_1 is anchored in F_1. Assume two S-points, v_1 and v_2, that are exactly alike outside of the F_1 parameter and that give different values to the output, $\forall x_1 A$, say T at v_1 and F at v_2. Because of the F at v_2, there must (by the semantics of the universal quantifier) be a point, v_4, that is exactly like v_2 outside of x_1, and that gives A the S_1-value F. Now define v_3 as follows: it is just like v_4 outside of F_1 (hence also just like v_4 on x_1), and just like v_1 on F_1. So v_3 is exactly like v_1 outside of x_1. Hence, since $\forall x_1 A$ is T at v_1, it must (by the semantics of the universal quantifier) be that A is T at v_3. But then v_3 and v_4 have the following features: They give different truth values to A, even though they are exactly alike outside of F_1. And this reduces to absurdity the claim of anchoring: Any change in F_1 that (all by itself) makes a difference to the truth values of the output also sometimes makes a difference (all by itself) to the truth values of the input.

Observe that even when we have put these examples in terms of sentences, it is really a pre-semantic fact that $\forall x_1$ is anchored in the domain parameter and not in the F_1 parameter. It is no mere accident of grammar.

Finally, note that for essentially 0-ary S-operations, being anchored (unanchored) comes to the same thing as leaving open (closing), which in turn is the same thing as the output S-meaning being open (closed).

5.2. Possible and impossible combinations of operation properties

There are four dichotomies. For compactness, permit us temporary use of unmemorable acronyms for the various properties of S-operations in relation to a set P of S-parameters.

DEFINITION 5.5. (Acronyms for dichotomies). Fix S and P.

0 vs. +: the dichotomy between essentially 0-ary and essentially +-ary.

C vs. O: the dichotomy between Closing and leaving Open.
T vs. L: the dichotomy between Translocal and Local.
A vs. U: the dichotomy between Anchored and Unanchored.

Then there appear to be sixteen combinations; but only eight are really possible. We deal first with the types that are possible, and then with those that are not. We close this section by giving special consideration to a combination that is possible but odd.

FACT 5.6. (Possible types of S-operations) *The following eight types of S-operations are possible in relation to a given set P of S-parameters:* 0CLU, 0OLA, +CLU, +CTU, +OLU, +OLA, +OTU, +OTA.

The most evident collapse occurs for essentially 0-ary S-operations; for these all that matters is whether the uniquely given output S-meaning is open or closed in relation to P.

PROOF. We give examples in terms of connectives taking a sentence A into some familiar sentence $\ldots A\ldots$, supposing whenever possible a familiar S_1-type semantics for these connectives. (We use $\{x_1\}$ for the set containing just the x_1 parameter.) In several cases phenomena of interest do not seem to appear in S_1; in these cases we appeal in rough terms to grammars and pre-semantic systems that are as familiar as possible.

0CLU examples. Two paradigms. (1) The 0-ary operation attached to the 0-ary grammatical function that produces the atomic sentence Fx_1 from the empty set of arguments; in the parameter-set $P_1 - \{F,\, x_1\}$. (2) The unary but essentially 0-ary operation taking any S_1-meaning into the constant T S_1-meaning; in the set P_1 of all S_1-parameters. This operation is attached, for example, to the unary grammatical function that takes any sentence A into the excluded middle, $(A \vee {\sim}A)$.

0OLA example. Paradigm: The 0-ary operation attached to the 0-ary grammatical function that produces the atomic sentence, Fx_1, from the empty set of arguments; in the parameter-set $\{x_1\}$. (This example was spelled out a bit under the definition, given above in Definition 5.4, of anchoring.)

+CLU example. This odd combination can be illustrated in identity theory, where it describes the operation attached to the connective that takes a sentence, A, into the sentence $\forall x_1 \forall x_2 (x_1 = x_2) \,\&\, A$; in, for example, the set of parameters $\{x_1\}$. See Fact 5.9 below for more information on the oddity involved. (We do not know if one can illustrate this oddity in quantification theory.)

+CTU example. Paradigm: The S_1-operation attached to universal quantification on x_1 (that is, the connective that takes A into $\forall x_1 A$); in the set $\{x_1\}$ of S_1-parameters.

+OLU example. Two paradigms. (1) The S_1-operation attached to the negation connective (that is, the connective that takes A into $\sim A$); in the set P_1 of all S_1-parameters. (2) The S_1-operation attached to universal quantification on x_1 (that is, the connective that takes A into $\forall x_1 A$); in the set $\{x_2\}$.

+OLA example. Paradigm: The S_1-operation attached to the universal quantifier; in the parameter-set {the domain parameter}.

Another paradigm can be found attached to the *Now:* connective introduced by Kamp into context-dependent tense logic. So suppose we have tense logic with both a time parameter and a time-of-context parameter. Consider the connective that takes A into *Now:A*, and the operation that attaches thereunto. In the time parameter this operation is merely +CTU. But in the time-of-context parameter, the *Now:* connective is +OLA, as we wished to illustrate.

Another example—albeit perhaps not paradigmatic—is also found in S_1: the S_1-operation attached to the connective that takes A into $A \& Fx_1$ is of type +OLA in the parameter-set $\{x_1\}$.

+OTU example. Seemingly odd but found in S_1: the S_1-operation attached to the connective that takes A into $(\forall x_1 A \vee \sim A)$; in the set $\{x_1\}$ of S_1-parameters.

+OTA example. Paradigm in tense logic: the operation attached to the future tense connective (the connective that takes A into *Will:A*; in the time parameter). Paradigm in modal logic: the *S4*-type operation attached to the necessity connective (the connective that takes A into $\Box A$); in the world parameter. This type is also found in S_1. The S_1-operation attached to the connective that takes A into $\exists x_1 A \& Fx_1$ is of type +OTA in the parameter-set $\{x_1\}$. ∎

Next we treat the eight impossible combinations.

FACT 5.7. (Impossible combinations) *The following combinations are impossible:* 0CLA, +CTA, +CLA, 0CTA, 0CTU, 0OTU, 0OTA, 0OLU.

PROOF. We show the following sub-combinations to be impossible: CA, 0T, and 0OU. The combination CA rules out the first four listed above, while 0T rules out the fourth through seventh. And finally, 0OU prohibits the last of those listed.

CA is impossible. It is trivial from the form of the definitions that closing implies unanchored.

0T is impossible. Fix S. Choose ψ and P. Suppose that ψ is essentially 0-ary, and use this to choose I_0 such that $\forall I[\psi I = I_0]$, and therefore $\forall I \forall v[\psi I v = I_0 v]$. So for arbitrary v, I_1, I_2, it must be that $\psi I_1 v = \psi I_2 v$. So, vacuously, ψ is local in P.

0OU is impossible. Fix S. We derive a contradiction from the supposal that ψ *(a)* is essentially 0-ary, *(b)* leaves P open, and *(c)* is unanchored in P. Choosing witnesses to *(b)*, let *(d)* $v_1 =_{\mathcal{P}_S - P} v_2$ and *(e)* $\psi I_1 v_1 \neq \psi I_1 v_2$. Choose I_2 so that *(f)* I_2 is a constant S-meaning, $\forall v_1 \forall v_2 (I_2 v_1 = I_2 v_2)$. Then *(g)* $\psi I_1 = \psi I_2$ by *(a)*, so that *(h)* $\psi I_2 v_1 \neq \psi I_2 v_2$ by *(e)* and *(g)*. Finally, *(d)* and *(h)* imply, via the "contrapositive" form of *(c)* given in 5.4, that $I_2 v_1 \neq I_2 v_2$. This contradicts *(f)*. ■

There is just one thread hanging. In the proof of Fact 5.6, we illustrated the possibility of +CLU in identity theory, but the example was peculiar. Here, to close this discussion, we offer a structural characterization of its oddness.

Definition 5.8. (Unique confinement). Fix S. An S-point, v_1, *uniquely confines* a set, P, of S-parameters iff every S-point that agrees with v_1 outside of P also agrees with v_1 on P: $\forall v_2[v_1 =_{\mathcal{P}_S - P} v_2 \rightarrow v_1 =_P v_2]$.

The oddity, if such there be, is that if v_1 uniquely confines P, then the variability that we should expect a set of S-parameters, P, to exhibit, even when we hold fixed the values of all other S-parameters, is missing when we start with the point v_1. The paradigm example of the unique confinement oddity is this. Almost always the x_1-parameter, in S_1, permits real variation no matter from which S-point, v_1, one starts. But consider a special case in which the domain parameter of v_1 is fixed at a domain with but a single member. Then, of course, there is only apparent variability to the x_1-parameter. Really, as long as we consider points v_2 that agree with v_1 outside of the x_1-parameter, and hence points, v_2, that agree with v_1 on the domain parameter, we shall obviously find that the x_1-value of v_2 is exactly the same as the x_1-value of v_1.

The following fact establishes that the combination +CL can occur only in the presence of unique confinement.

Fact 5.9. (The combination +CL—Essentially +-ary, Closing and Local—implies unique confinement) *Given S, the* +CL *combination of essentially +-ary with closing and local on P is impossible unless S sometimes permits unique confinement of P. In other words, if no S-point uniquely confines P and if ψ is both local in and closes P, then ψ is essentially* 0*-ary.*

PROOF. Fix S. Suppose that *(a)* no S-point uniquely confines P, and that *(b)* ψ is local in P. Further suppose that *(c)* ψ closes P. We show that *(z)* ψ is essentially 0-ary. Choose I_1 and I_2; we need to show *(y)* $\psi I_1 = \psi I_2$, which is to say, for arbitrary v_0, that *(x)* $\psi I_1 v_0 = \psi I_2 v_0$.

Define the S-meaning, I_3, by giving its value for each S-point, v, by cases as follows: *(d₁)* if not ($v =_P v_0$) then $I_3 v = I_1 v$, and *(d₂)* if $v =_P v_0$ then $I_3 v = I_2 v$. We can obtain *(x)* immediately by showing that *(w₁)* $\psi I_2 v_0 = \psi I_3 v_0$ and that *(w₂)* $\psi I_1 v_0 = \psi I_3 v_0$.

To show *(w₁)* we appeal to locality. By *(d₂)* we have *(e)* $\forall v_2 [v_0 =_P v_2 \rightarrow I_2 v_2 = I_3 v_2]$. Then we obtain the desired *(w₁)* from *(e)* by locality, *(b)*.

The argument for *(w₂)* begins with an appeal to *(a)* the absence of unique confinement of P by, in particular, v_0: There is a v_2 such that *(f₁)* $v_0 =_{\mathcal{P}_S - P} v_2$ but *(f₂)* not $v_0 =_P v_2$. Choose v, and suppose that *(g)* $v =_P v_2$. Combining *(g)* with *(f₂)* gives the falsity of ($v =_P v_0$), which with *(d₁)* gives $I_1 v = I_3 v$ under the hypothesis *(g)*. So *(h)* $\psi I_1 v_2 = \psi I_3 v_2$ from locality *(b)*, in analogy to the argument for *(w₁)*. Also *(j₁)* $\psi I_3 v_2 = \psi I_3 v_0$ and *(j₂)* $\psi I_1 v_2 = \psi I_1 v_0$ by *(f₁)* and the fact *(c)* that ψ closes P. So by a chain of identities from *(h)* and *(j₁)* and *(j₂)*, we have *(w₂)* as desired, which completes the argument. ■

We close by recording without proof an obvious connection, for essentially 0-ary S-operations, between types of S-meanings as given in Definition 4.1 and types of S-operations as discussed in the present section.

FACT 5.10. (Properties of essentially 0-ary S-operations) *Fix S. Suppose that ψ is an essentially 0-ary S-operation (Definition 5.1), and that its unique output is the S-meaning, I. Then (1) ψ is absolutely local, i.e., local in the set of all S-parameters. (2) Whether ψ leaves P open or closes it depends entirely on whether I is open (dependent on) or closed in (independent of) P. (3) Whether ψ is anchored or unanchored in P depends entirely on whether I is open (dependent on) or closed in (independent of) P.*

6. Application to other pre-semantic systems

Suppose we take a modal logic for a language with a necessity connective and a "it's true in reality that" connective. For such a logic, let us construct a very standard pre-semantic system, S_2. The S_2 parameters are these: set of worlds, relation of relative possibility, real world, world of evaluation, proposition letters. When we flatten, it is obvious that truth of a sentence in many a standard modal logic is relative to each of these parameters. So the S_2-values are still the two truth values.

Fixing attention on a single S_2 point, v, auxiliary values must be as follows. *Of the set-of-worlds* parameter at v: any set. *Of the relative-possibility* parameter at v: a binary relation on the value of the set-of-worlds parameter at v. *Of the real-world* parameter at v: a member of the value of the set-of-worlds parameter at v. *Of the world-of-evaluation* parameter at v: a member of the value of the set-of-worlds parameter at v. *Of each proposition letter* parameter, a function from the value of the set-of-worlds parameter into {T, F}, or, equivalently, a subset of the value of the set-of-worlds parameter at v.

These standard Kripke-style choices give us the four primitives: S_2-parameters, S_2-points, S_2-values, and S_2-auxiliary values. Then S_2-meanings and S_2-operations are defined uniformly from these, the S_2-meanings being functions from S_2-points to S_2-values, and the S_2-operations being mappings from S_2-meanings into S_2-meanings. Thus we have specified all six ideas needed for a pre-semantic system.

Consider the S_2-operations, ν and ρ, associated respectively with the necessity connective and the "it's true in reality" connective. Imagine that ν reflects a Kripke "relative possibility" semantics for the necessity connective and that ρ encodes the usual modal semantics for the "it's true in reality" connective. Also let q stand, in context, both for the proposition letter q qua categorematic expression and qua parameter, and let I_q be the S_2-meaning associated with q qua categorematic expression.

- $I_q v = (q_v)(\text{world-of-evaluation}_v)$.
- $(\nu I)v = \text{T}$ if for all z, if ($z \in \text{set-of-worlds}_v$ and $\langle\text{world-of-evaluation}_v, z\rangle \in \text{relative-possibility}_v$) then $I([z \;/\; \text{world-of-evaluation}]v) = \text{T}$; and otherwise $(\nu I)v = \text{F}$.

  That is, νI is true at v iff I is true at all points just like v, except that the world of evaluation has been shifted to a member of set-of-worlds$_v$ that is relatively-possible$_v$ at world-of-evaluation$_v$.
- $(\rho I)v = I([\text{real-world}_v \;/\; \text{world-of-evaluation}]v)$.

  That is, ρI is true at v iff I is true at the point that is just like v, except that the world of evaluation has been shifted to the real world of v.

Our notation is evidently difficult to read and write; were we to plan on using these concepts much, abbreviating definitions would be in order. The only point we wish to make, however, is that these operations have the following relations to the various parameters of S_2. We use the notation of Definition

5.5. In this table we put the parameters down the side (including parameters for proposition letters q and r) and the operations across the top, recalling for the latter that we are using ν for necessity and ρ for reality, and using q for the 0-ary operation producing the proposition letter q as a categorematic expression.

	ν	ρ	q
set of worlds	+OLA	+OLU	0CLU
relative possibility	+OLA	+OLU	0CLU
real world	+OLU	+OLA	0CLU
world of evaluation	+OTA	+CTU	0OLA
proposition letter q	+OLU	+OLU	0OLA
proposition letter r	+OLU	+OLU	0CLU

Most of this typing is revelatory, for example, the confusingly different relations of the necessity and the reality operations to the real world and to the world of evaluation. Observe that with a "natural" S5 semantics, necessity comes out +CTU in the world of evaluation, instead of +OTA. One should therefore be slow in making a bald statement such as "in modal logic, necessity . . . ."

Observe that when a proposition letter is taken as a parameter, its *auxiliary* S_2-value at a point is an "intension" (function from worlds into truth values). Nevertheless, when the same letter is taken as a categorematic expression (a sentence), its S_2-value at that same point is a truth value. This is just right, and exactly in accord with both common practice and Carnap's method of extension and intension. Remark also that if one uses proposition letters as bound variables in modal contexts, one treats such a variable in exactly the same way as the constants: Its S_2-auxiliary value at a point is an intension, whereas its S_2-value at a point is a truth value. There is no difference in the underlying semantic treatment of propositional constants and propositional variables. Carnap's logical insight, amusingly but unfortunately spoofed on metaphysical grounds by Quine (in the appendix to [3]), was to see that the same should be true of "atomic" symbols of every type. Predicate symbols, whether constant or variable, should be awarded intensions at points as S_2-parameters, and extensions (of the proper type) at points if taken as categorematic.[11] Standard modal logic agrees.

[11] [2] took the logical policy of uniformity even further, with results that have, alas, been almost entirely ignored by the community of modal logicians. Take a higher-order predicate such as "is contingent." It is obvious that application of this predicate to a first-level predicate, F, must take into account the entire "intension" of F, not just its "extension" at the point of evaluation. In other and clearer words, predication at the

What looks odd and even wrong-headed, however, from the point of view of flat pre-semantics, is that standard modal logic treats individual variables (always) and individual constants (sometimes) differently from proposition letters and predicate letters. Standard modal logic forces the S_2-auxiliary values at points of individual variables (always) and individual constants (sometimes) to be simple individuals instead of intensions (so-called individual concepts). This is metaphysics, not logic. Flat pre-semantics makes it plausible to distinguish S_2-auxiliary values at points from S_2-values at points. Then one sees the logical wisdom of treating individual symbols in exactly the same way as proposition letters for predicate letters: intensions for S_2-auxiliary values and individuals (extensions) for S_2-values. Absent metaphysical ideology, it figures: Atomic symbols of every type, whether constant or variable, and of whatever type, should—if one is guided by logic—be treated alike.

One last point about modal logic. We observed in the discussion of S_1 that identity, like truth functions, is absolutely local (local in every parameter). This is a *logical* remark, and supports the following: When an identity is added to a Carnap-Bressan type of modal logic with individual concepts, it is *contingent* identity rather than strict identity which is absolutely local, and which is therefore the proper logical descendant of identity as used in nonmodal contexts. Only an intrusion from some particular metaphysics would lead one to think differently.

The more complicated the language, the more complicated the pre-semantics, and the more helpful concept-sorting can be. Consider the logic of indeterminism as most simply represented in branching time.[12] One has concrete momentary events called *moments* that are arranged in a tree, and one has maximal chains in the tree that are called *histories*. One never gets straight on indeterminism unless one becomes aware that truth needs to be parameterized by both moments and histories. Having become clear to this extent, it helps enormously to avoid foolish "logical fatalism" arguments if one realizes that all ordinary tense and temporal constructions, although translocal in the moment parameter, are *local in the history parameter*. Therefore they work exactly the same as they do in ordinary linear

---

higher order should be translocal in the world-of-evaluation parameter. By uniformity, first order predication should also be translocal, taking into account the S_2-value of the terms to which the predicate is applied when the world-of-evaluation is varied. Bressan provides significant illumination by putting this uniformity firmly into effect. (Naturally, having made room for translocal predication in the pre-semantics, one will find that "most" predicates turn out to be local in the world of evaluation.)

[12] See for instance chapter 8 of [1].

tense logic. This is a simple logical point, but one that seems difficult to keep in mind. Having a *word* for it may help.

References

[1] Belnap, N., M. Perloff, and M. Xu, *Facing the future: Agents and their choices in our indeterminist world*, Oxford University Press, Oxford, 2001.

[2] Bressan, P., *A general interpreted modal calculus*, Yale University Press, New Haven, 1972.

[3] Carnap, R., *Meaning and necessity: a study in semantics and modal logic*, University of Chicago Press, Chicago, 1947.

[4] Kaplan, D., 'Demonstratives: an essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals; and Afterthoughts', *in* J. Almog, J. Perry and H. Wettstein, eds., *Themes from Kaplan*, Oxford University Press, Oxford, 1989, pp. 481–563; 565–614.

[5] Montague, R., 'Pragmatics', *in* R. Klibansky, ed., *Contemporary philosophy: a survey*, Florence, 1968, pp. 102–122.

[6] Tarski, A., 'The concept of truth in formalized languages', *in Logic, semantics, metamathematics*, The Clarendon Press, London and Oxford, 1956, pp. 152–278. Presented to the Warsaw Scientific Society in 1931.

[7] van Benthem, J., *Manual of intensional logic*, CSLI Publications, Stanford, 1988. Second edition 1991.

Nuel Belnap
Department of Philosophy
University of Pittsburgh
Fifth and Bigelow
Pittsburgh, PA 15260
belnap@pitt.edu