

Truth is dead; long live the truth. Commentary on *Conjoining Meanings* by Paul Pietroski

Gillian Ramchand 

Department of Language and Culture, UiT
The Arctic University of Norway, Tromsø,
Norway

Correspondence

Gillian Ramchand, Department of
Language and Culture, UiT The Arctic
University of Norway, Breivika N 9037,
Tromsø, Norway.
Email: gillian.ramchand@uit.no

Pietroski successfully dismantles the idea of a formal semantic theory based on direct truth conditions and offers new and formally constrained alternatives. In this paper, I summarize the arguments but also provide a number of test cases to show that refusing to accept Pietroski's conclusions condemns the field to constantly restating and technically evading its own self-created paradoxes. In the final section, I offer some positive proposals in the spirit of the Pietroskian enterprise with respect to thematic roles.

KEYWORDS

concepts, extensions, imperfective paradox, modals, thematic relations, truth conditions

1 | TRUTH IS DEAD

It is easy to be misled by the measured quality and respectfulness with which Paul Pietroski deals with his subject matter into thinking that it is not what it in fact is—a radical and provocative defense of the claim that truth-based theories of meaning are inadequate to the task, and should now be abandoned, that is, the claim in D is false.

- (1) D. If L is a Slang, then some theory of truth for L is the core of a correct theory of meaning for L.

If you are a working formal semanticist of natural language and are not shocked and/or exhilarated by the content of this book, then in my opinion you were not reading carefully enough.

In this first section, I will summarize what I take to be the main claims and deep reasons behind the rejection of semantics with truth values. The main idea is that even though we clearly want to understand the use of language to “say things about the world,” that does not mean that natural

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2019 The Author. *Mind & Language* published by John Wiley & Sons Ltd.

language “meanings”—by which we mean the other side of the memorized link that connects pronunciations to what they invoke—must literally be extensions. According to Pietroski, while we sometimes use meanings “to build truth evaluable thoughts” or “refer to entities in the world,” the meanings of lexical items themselves are the cognitive precursors of those acts. By starting from the assumption that sentences denote truth or falsity and that meanings of lexical items are functions from entities to truth values, or some recursive type associated with those extensional primitives, one effectively begs the question. In other words, it tells us nothing about what precisely the cognitive ingredients are and how human minds combine them to create these complex truth evaluable expressive acts.

1.1 | Problems too hard for Frege-Tarski-Church toolbox

There are a certain number of classical paradoxes for the current system that are well known. The first is the fact that even with the addition of possible worlds (Lewis, 1986), truth conditional meanings do not seem to be quite fine-grained enough to do justice to Fregean “senses.” The problem arises with logically equivalent statements, which should give the same truth value in any possible world, but which are obviously distinguishable in terms of sense or naive meaning (chapter 1).

$$(2) \quad 2 + 2 = 4 \neq 150 + 160 = 310$$

Problems for truth theories also arise in interaction with the use of the predicate truth itself, which leads to the Liar's Paradox (Chapter 4). These problems are well known and there has always been a tacit agreement among working semanticists at any rate to just ignore them and get on with things.

There is also the problem of semanticity. Consider the truth theorems that Davidson urges on us as things that a suitable theory of meaning should be able to derive.

$$(3) \quad \text{“Snow is white” is True if and only if } snow \text{ is white.}$$

This is all very well and good, but it is not just important for the theory to derive (3), it also must *not* generate (4).

$$(4) \quad \text{“Snow is white” is True if and only if } snow \text{ is white and } 2 + 2 = 4$$

It is hard to see how an extensional theory of meaning based on truth conditions avoids this, without resorting to a condition on the use of *particular lexical ingredients* (Chapter 1).

This leads us inexorably to the final problem, and the real deal breaker in my opinion. This is the problem of the meanings of open class lexical items—their flexibility, their polysemanticity, even while exhibiting productive and generative composition with each other. One could argue that building truth theorems of the type in (3) is already hard, and good enough, and that it really is not the immediate job of the formal semanticist to elucidate the meanings of the individual lexical concepts *snow* and *white*. (I have indeed myself justified the Davidsonian enterprise in precisely these terms to generations of first-year semantics students.) However, it simply will not do. The problem is that the meanings of open class lexical items are not monolithic blobs of coherent conceptual stuff encapsulated and insulated away from the combinatorial system, the fact is that (a) they are conceptually polysemous and

(b) they undergo productive compositional processes with each other. The latter point shows that understanding their behavior is an important component of understanding the central properties of the human language system and its powers of productive meaning generation. So, we *cannot* ignore the word “book”, to take an example, and we *cannot* ignore the fact that “book” in (5-a) refers to the content, while in (5-b) it refers to the physical artifact. The by now standard approach from formal compositional semantics is to reverse engineer the meaning required based on the truth conditions for the sentence as a whole. This would, effectively, create two distinct lexical meanings for (a) and (b), with no real engagement with the systematicity of the alternation, or the fact that it occurs for free in most human languages that I am aware of (i.e., not mediated by any kind of derivational morphology, or distinct lexicalization). In (5-d), natural language seems to allow it to refer to both things simultaneously, throwing up insurmountable problems for standard denotations of type $\langle et \rangle$ (see the computational literature on dot types for a discussion of the hardness of this problem).

- (5) (a) John loves to read books.
 (b) John has no room for all his books.
 (c) This book was written by Pietroski.
 (d) This interesting book fell off the shelf and broke my toe.

This is not a problem that standard formal semantics normally concerns itself with, which means that it is abdicating half of its job.

1.2 | Problems too easy for the Frege-Tarski-Church toolbox

The problem here essentially is that once you have type theory and abstraction, then it is possible to create higher types to get you out of any jam that compositionality might force upon you. Consider the standard problem of the denotation of adjectives.¹ In postcopular position, one wants to give the adjective a semantic type $\langle et \rangle$. So, in the case of (6-a), the denotation of “big” is a function that takes entities in the domain and maps to TRUE if and only if that entity is in fact big. In a parallel way, we can think of the noun “elephant” as denoting the property of elephanthood, or more precisely, that it is the function that delivers TRUE if and only if the thing it applies to is in fact an elephant.

- (6) (a) Dumbo is big.
 (b) Dumbo is an elephant.

Let us put aside for the moment that we need to include some sort of contextual restriction on the standards of comparison for bigness here, something that humans seamlessly accomplish, but which the formal system could in principle accomplish too with suitably inserted contextual variables. Dumbo might be big for an elephant, but he certainly is not big for a planet. For now, let us concentrate purely on the type. If we place the adjective now in a position where it modifies a noun, we get a sentence like (7).

- (7) Dumbo is a big elephant.

¹ This discussion is deliberately chosen to echo the discussion in Heim & Kratzer 1998 concerning predicate modification.

Here we need to combine “big” with “elephant” in such a way that the result is something of $\langle et \rangle$ that is true of “big elephants.” So what do we do here? If we keep to the principle that all meaning is represented by extensional functions, and all meaning composition is function-argument combination, then we need to say that in this case the word “big” is something of type $\langle \langle et \rangle, \langle et \rangle \rangle$. In other words, it combines with a one-place predicate to build a new one-place predicate. The new one-place predicate is characterized by the fact that it requires the two original one-place predicates to be separately true, with the caveat that the interpretation of the one is subtly contextually modulated by the other, another contextual variable of the sort required by (6-a) anyway should do the trick here. But there is an alternative. Heim and Kratzer, to their credit, give us one. The alternative is to say that both “big” and “elephant” are of type $\langle et \rangle$ and that when two such one-place predicates get together, there is a special mode of combination that natural languages allow whereby two monadic predicates join to form another monadic predicate, true of both separately. (This by the way is exactly the proposed M-join operation that Pietroski claims is the fundamental primitive building block of human language combination, and which together with D-join, entirely replaces the function-argument strategy.) What Heim and Kratzer (H&K) have to say here is instructive. They say that *it does not matter*—either way will work, and it depends on where you wish to amass your costs. The real question now is whether this is a very meaningful choice or a meaningless one. The implicit position of H&K seems to be that the debate is ultimately not very interesting once we get the truth conditions to compose correctly. But the exercise highlights an important property of the system, namely that having functions of higher types means that you are never forced to make substantive linguistic claims about what sorts of combinations human minds traffic in. In this particular case, the Fregean move seems quite benign and does not fare too badly in the matchup with predicate modification. But the problem scales up, and the Fregean types get much more complicated, while if Pietroski is right, monadic (and a highly restricted form of dyadic) predication exhausts what we will eventually need.

Pietroski points out in his book that the Frege-Tarski-Church playbook (having been designed ultimately for rather different purposes) offers the possibility of many more types of function than are ever used in the analysis of human meaning combinations (p. 152), and even those that are used could be rethought in more minimal terms. So, it seems that it is the job of linguists to actually make decisions of the type described above. Far from it being an innocent question of technical machinery involving aesthetic judgments of simplicity, it is precisely the job of linguistic semanticists to classify the primitive cognitively realistic meaning units and the substantive types of meaning composition we actually find. This is the positive part of the work, and will operate as a theory of the substantive constraints on the power of the lambda calculus. So for every detailed problem of meaning composition natural language throws up for our consideration, and there are many, the solution offered by lambda calculus is simply *too easy, and basically trivial to implement*,² in the absence of a real theory of meanings and meaning combination.

1.3 | The Pietroski proposal

If you are going to argue against the most powerful tradition of natural language semantics, you had better spent over 300 pages justifying yourself and go over the arguments very rigorously. In the first part of the book, Pietroski carefully deconstructs the intellectual history of the Fregean tradition, pointing out the formal stages on the journey, and explicating the goals of the endeavor. On the way, he shows that the final position arrived at is neither inevitable nor even the obviously correct tool for natural language semantics. As a linguist with training in formal semantics, one does not standardly

² It is also too complex, in the sense of creating a complexity of symbology that does not match the relative simplicity of the linguistic primes.

get this history presented in such a critical and deconstructed way. Rather, one gets a history of progress, which teleologically leads to a sanctioned final state of the art. One gets the impression that the initial state of knowledge was like an incredibly complex chess board position, which could and was in the end converted to a win for White. We as individuals encountering the initial position need to be led systematically through the steps that got us to that goal, a goal which had to be arrived at through many extremely smart scholars successively contributing elements of the solution. Like a computer assessment of the correct line to the inevitable solution—the final goal is never the question, how human minds eventually got there is the exciting question. But this all presupposes that the solution was the only and correct one. Pietroski shows it for what it actually was, a fascinating journey of human intellectual development and theory construction, involving particular human minds, with historically contingent decisions along the way (see the work of Thomas Kuhn). This is an important part of the book (Chapters 1–3), and is necessary reading for allowing the reader to begin to step outside of the standard narrative and at least have an open mind to the rather different solutions that Pietroski in the end proposes.

Deconstruction is never enough. I am sure that many scholars have felt and expressed discontent with the dominant paradigm, and none of the complaints and reassessments are completely unique to Pietroski, although they are brought together here to powerful effect. The real virtue of the book is the fact that it is coupled with a positive contribution—the presentation of a set of proposals and ideas that could replace our older conceptions of meaning. To jump ship, we need to be shown that there are viable alternatives that have a chance of meeting the desiderata and standards already met by the older tradition, and show how those alternatives have the potential to go even further. Pietroski's own proposals are interleaved throughout the book, to greater and greater degrees of explicitness as the book progresses. The critical deconstructive voice alternates with the positive and inquiring voice throughout, the former dominating the latter somewhat in the early chapters and the latter providing the melody line almost entirely in Chapters 6 and 7.

In the following list, I summarize in a brief and crude way the main substantive proposals of the book. There are many interesting concrete proposals about particular phenomena along the way that I will not be able to comment specifically on. In particular, the treatment of plurality using George Boolos's idea for the interpretation of plural variables is original and has far-reaching implications, which deserve an article on its own. The treatment of quantification, similarly, is one which is explicit and novel, and sufficiently different from standard treatments that it deserves a full discussion in its own article. I will refrain from engaging in these two domains in this commentary, reserving my most specific substantive comments for the area of event structure and thematic roles which is closest to my own area of expertise.

- A. Meanings are neither concepts nor extensions; they are instructions to “fetch” (in Pietroski's terms) a potentially polysemous network of concepts
- B. Basic meanings take the form of one of only two basic sorts, regardless of the complexity of the concepts they evoke: *monadic predications* (M Predicates), and certain restricted but effective *dyadic relations in the lexical domain* (D Predicates) allowing certain amount of extra combinatoric and expressive power.
- C. Meanings combine in ways that are severely restricted compared to the full combinatorical power of the type-based approach using functions. These are defined so as to ensure that the meaning of two joined meanings properly includes those individual pieces.³ Pietroski defines

³ This is a stronger notion of compositionality than simply saying that the meaning of a complex form can “be expressed in terms of” the meanings of the parts as in the Fregean system.

two basic semantic combinatoric operations: M Predicates combine with M Predicates by M junction and M predicates combine with D predicates by D junction.

- D. Events are necessary for the understanding of verbal meanings, but they give rise to paradoxes when coupled with a truth-based semantics. Verbs should be seen as ways of framing events in the world. Thematic roles are generalized dyadic predicates that allow the combination of verbs with their arguments in systematic ways. (Chapter 5)
- E. Tense, negation, and quantification can all be understood in these more restricted terms, and do not require the higher types often proposed. (Chapter 7)

But first, if truth-based theories of meaning are dead, and we are to internalize and accept this reality, then there are a couple of stages we need to go through first.

2 | LOSING A LOVED ONE: THE SEVEN STAGES

2.1 | Denial

One reaction to all of this is to deny that there is a problem. After all, formal semantics is a huge success story and brilliant analyses of natural language phenomena are continually being produced, and problems are being solved. The problems that Pietroski raise are based on a misunderstanding of the different levels of analysis well known from Marr (1982), and the place of formal semantics based on truth within that classification. Essentially, these theories of meaning are clearly at the computational level in Marr's typology, and should not be asked to do representational and algorithmic duty. And while it would be no doubt nice to have a theory of representation and algorithm, we are very far from being at that stage now in our current understanding of the mind/brain. After all, Marr himself argued that a good solid theory of the computational level is necessary if one is to (eventually) move to the next step.

It is true that there are very many good papers written in this framework that really make progress on generalizations about language meaning. I deconstruct one representative analysis here, as an example of excellent research that is explicitly engaged in tackling linguistic questions about how meaning formatives are put together, but which is actively impeded by the received framework.⁴

2.1.1 | Modals and past tense

Condoravdi (2002) deals with the interaction between modals and temporal interpretation. Traditional accounts simply point out that *John might go to the party* does not have a past tense reading but a future oriented one, and therefore that modals are not tensed in any straightforward way, even though *might* seems to bear old past tense morphology. Condoravdi points out that in order to understand the interaction of modals with tense inflection one needs to distinguish between the temporal perspective of a modal and its temporal orientation. The *perspectival situation* is the situation at which the *potential* for the prejacent event is asserted, and the *at-issue situation* is the situation expressed by the prejacent event description. In (8), we can see that the *at-issue* situation

⁴ The moral of this part of the story eventually is going to be that even in the best research we can detect the shortcoming of the starting assumptions for pushing the accounts further.

can be projected forward from the perspectival situation as in (8-a), be at the same time as the perspectival situation as in (8-b), or be earlier than the perspectival situation as in (8-c).

- (8) (a) John might go to the party.
 (b) John might be in his office.
 (c) John might have won the race.

One way to capture this is to specify the semantics for MIGHT directly in two versions, one forward shifting and the other non-shifting with respect to the possible worlds considered. One might even stipulate a third composite modal MIGHT-HAVE which requires the possible worlds to precede the utterance time. Condoravdi (2002) gives the following three possible denotations corresponding to such a view. In the following, MB designates the modal base that a modal depends on for its interpretation (Kratzer, 1977).

- (9) (a) Forward shifting modals:
 $\text{MIGHT}_{MB}^1 \phi$ is true at $\langle w, t \rangle$ iff there exist w', t' such that w' in $\text{MB}(w, t)$, $t < t'$ and ϕ is true at $\langle w', t' \rangle$.
 (b) Non-shifting modals:
 $\text{MIGHT}_{MB}^2 \phi$ is true at $\langle w, t \rangle$ iff there is w' in $\text{MB}(w, t)$, such that ϕ is true at $\langle w', t \rangle$.
 (c) Backward-shifting modals:
 $\text{MIGHT-HAVE}_{MB}^1 \phi$ is true at $\langle w, t \rangle$ iff there exist w', t' such that w' in $\text{MB}(w, t)$, $t' < t$ and ϕ is true at $\langle w', t' \rangle$.

Quite rightly, Condoravdi wants to do better than this, and with respect to the past *at-issue* situation in [c], the reading seems to be related to the presence of the *have* auxiliary. In the case of forward shifting versus cotemporaneity of the *at-issue* situation, it turns out that the interpretation here is closely entwined with the epistemic vs. circumstantial distinction, as well as with the stative versus eventive distinction. I do not go into the details here, but refer the reader to Condoravdi (2002) for a clear and essentially correct description of the intricate linguistic generalizations she uncovers in this domain. However, I want to focus on one particular subcase, which is the one where the *at-issue* situation is in the past of the *perspectival situation*. This reading is crucially dependent on the presence of the perfect auxiliary *have* in the prejacent. In addition, sentences like (10), are in fact ambiguous, having an epistemic uncertainty reading (a) in addition to a counterfactual one (b).

- (10) Mary could have won the election.
 (a) Epistemic uncertainty: "...Let's go check!"
 (b) Counterfactual: "...If she had only campaigned more"

Condoravdi (2002) eventually proposes an LF movement solution to account for the two different readings. The first assumption is that the perfect auxiliary *have* actually introduces a relative PAST operator which interacts with the modal operator. Since the modal operator and the tense operator operate over propositions, it is possible to take them in principle in either order. Under this view, it seems elegant to notice that the epistemic reading looks like it can be derived from the epistemic possibility modal scoping over the PAST operator, and the counterfactual reading as being closer to the past operator scoping over the epistemic possibility (with a little help from pragmatic strengthening).

However, what looks like an elegant and satisfying solution simply does not work once one looks closer at the facts. Basically, this is because as a free scopal operation, it massively overgenerates. It overgenerates fatally at the local level for English, since the counterfactual reading is simply absent for modals like *must* and *may*, as in (11).

- (11) (a) Mary must have won the election.
 (b) Mary may have won the election.

And indeed for any modal in English, that does not bear moribund tense morphology. This indicates that it is the modal itself that optionally bears the possibility of PAST here— not a scopal interaction with *have*, but if anything with its own moribund past tense morphology. This now raises the question of why *have* is necessary in the construction in the first place. The solution also overgenerates fatally at the level of typology because if such scopal inversions of tense and circumstantial modality were always possible, why do we always find $T > \text{Modal}_{\text{circ}}$ as a typological generalization?

The problem with this excellent paper is the assumption that is built into the toolbox being used, namely that the modal operator takes a function from sets of possible worlds to truth values. This means that we cannot make a distinction between modals or between modals and tense with respect to what they combine with semantically, which in turn makes it impossible to connect this in any way to syntactic height. In the case of the Condoravdi paper, syntactic ordering (in this case of the modal and the *have* auxiliary) is deliberately ignored, or rather undone by a semantic operation to get the interpretational effects. But further linguistic facts show this to simply fail as a description of the facts. Not surprisingly, I would argue, since in the building up of linguistic meaning, *order does matter*.

2.2 | Pain, anger and bargaining

So, order does matter, and the toolbox we have seems to flatten everything out into functions with extensional primitives as their inputs and truth values as their final output. This means that we have no way of modeling the kind of “partial content” that I think actual linguistic forms traffic in. And as Pietroski points out, for natural language, there is no logical necessity for our truth evaluability to be so direct. Rather, we need to see how sentences constructed systematically by natural language, eventually end up building something that *is* truth evaluable, but this does not mean that meanings themselves, the linguistic ingredients of the sentence, are extensions.

There are a number of ways to bargain one's way out of this position even if one accepts the logic of Pietroski's arguments. One obvious way is to make a kind of bargain, or pact, whereby we can continue to use the classical system, but promise to use it responsibly with full understanding of its potential limitations. The reason for making the bargain would be the idea that even though we do of course eventually need to make the connection to actual linguistic partial contents, such a move is premature given our current understanding of the human mind/brain. The reason that this plea bargain does not work is that even at the computational level, this insistence on meanings being extensions very quickly hits a brick wall with respect to the genuine generalizations it can deliver, and gives rise to the same problem over and over again. In the worst case, we create a mini industry, attempting to solve by increasingly more baroque and technical means a problem that is a direct artifact of our own flawed technical starting point. My example of this situation is the imperfective paradox which I will describe next.

2.2.1 | The imperfective paradox

This problem goes back to the very beginning of Montague semantics, and involves the difficulties that arise when one attempts to express the truth conditions of the English progressive in terms of the truth conditions of the verb it is based on.

- (12) (a) John crossed the street.
 (b) John is crossing the street.

Bennett and Partee (1978) noticed that you could not give a semantics of the progressive that said it was true at a proper subinterval of the interval at which the nonprogressivized sentence is true. This is because if you say that, you commit yourself to John actually getting across the street at some future time. To get out of this unwanted completion entailment, one needs to say that the complete interval corresponding to “John cross the street” is not true in the real world necessarily, but in some possible world, plausible given the starting conditions (i.e., “inertial” in Dowty's, 1979 terms), leading out from the situation that the progressive sentence describes (Dowty, 1979). In fact, this is also not good enough because even if all inertial worlds would never lead to a successful street crossing (because, for example, a large truck is rushing toward John without him realizing it), we can *still* say that “John is crossing the street”. So, inertial worlds are a little too large and crude. What we need is a notion of continuation that relativized to the event and modulated by event internal properties such as intention and abilities of the agent (Landman, 1992). And analyses of the imperfective paradox continue to this day, with more and more fine-grained scenarios forcing the semanticist to precisify the nature of the modal base, and/or call into question the nature of the quantification involved. It is fair to say that we have had nearly 40 years of possible worlds accounts of the humble progressive. The overwhelming consensus of the field is that some such analysis is the correct one for this construction, and for imperfectives in other languages.⁵

However, all of the possible worlds accounts still fall short of complete objective explicitness when it comes to characterizing the progressive construction semantically in terms of the corresponding non-progressive. In all cases, the appeal to possible worlds still leaves an unexplained residue completely independent of the possible worlds mechanisms themselves. In the case of Landman, it is his appeal to the “stage-of” relation, in Portner (1998) it is the relativization to event descriptions, in Hallman (2009) situational version it is the relation R “the relevant subpart relation”. The essential question of “What does it mean to be an in-progress version of an event?” remains a primitive. The problem here is not that we need to express the progressive in terms that *build on* the nonprogressivized version—that is just a basic compositional desideratum. The problem is that the toolbox forces us to start with an extensional meaning for the nonprogressivized version. But the essential meaning of a verb (unprogressivized) cannot be captured via extensions, in Pietroski's view, so this strategy is doomed to failure. And the progressive is intuitively related to the main verb in a way that does not require truth. The progressive is surely built from the simple verb via some systematic meaning operation, but one in which meanings are not extensions.⁶

There is in fact good evidence that the building of the derived stative participle in *-ing* is in fact cognitively basic. It has been known for a long time that the progressive participle in *-ing* is one of

⁵ The chief rival account, Parsons, 1990's nonmodal account in terms of In-Progress state, seems to be rather unfashionable, although to my mind its difficulties are no greater than the ones that plague the modal accounts, and its benefits are substantial.

⁶ See Ramchand, 2018 for such an account.

the very earliest pieces of morphology acquired by English children. It is acquired between the ages of 19–28 months, and appears *before* both irregular past tense (which in turn appears often before regular past inflection) and the copula (Brown, 1973; Owens, 2001). The use of the *-ing* participle thus appears before any actual tense inflection or modal expression, and is used correctly immediately.⁷ A fully modal and intensionalized analysis of the progressive would require us to believe that English children acquire a modalized meaning accurately before they are 2 years old, and always do so before they even have the ability to express tense or use modal auxiliaries. The pragmatic complexity of inferences connected to the setting up of modal bases and ordering sources is supposed to be something that children need some social and interactional maturity in order to develop. But standard accounts seem to assume that they can do this even before they pass theory of mind tests. The point of formal explicitness is so that we do not hide from ourselves our unexplained assumptions. But here the unexplained part is the core of the progressive meaning itself—the idea of what it means to be a subpart of a particular event described in a particular way.

So no, we cannot just continue to use the old system and promise to use it carefully. All the puzzles and paradoxes are the same, all the dead ends have the same deep problem at their heart. All the technical solutions contain the same self-defeating trick (structure the possible world space in an increasingly elaborate, context-sensitive way, often depending on *the actual words used* in building the construction). This shows that we really have not understood the meanings of words in a way that is helping us to solve the compositional issues.

The bottom line here is that we can grant that formal semantics is a computational theory (in the sense of Marr (1982)) of meaning and still hold it to certain standards of correctness. Even if successful in its own terms, correctness needs to be understood at least partly in terms of a theory's potential to inform and constrain theories of representation and algorithm. The current theory, even when wielded by extremely intelligent and knowledgeable linguists ends up forcing us into dead ends and paradoxes. All of which could be avoided by dropping the empirically undermotivated assumption that meanings are immediately extensional.

2.3 | Sadness followed by upturn

I suspect that there will be many who read Paul Pietroski's new book who do not get this far in the process. It is a fact about intellectual paradigms that they are remarkably slow to change, and individuals who have invested much research time within a particular paradigm are probably impossible to convince. I think Pietroski has done a compelling job of it, but I wonder how many working semanticians out there who were not already convinced, would be swayed by his arguments. For those who are convinced that a sea change is needed, sadness at losing an old toolbox can quickly lead to the excitement that comes with new possibilities.

And it is not so bad actually to have had 50 years of this particular tradition of formal semantic theorizing. There is a lot we know because of it that we did not know before. In concentrating on the externally verifiable truth conditions of sentences, we have a sort of E-semantics descriptively in place, that can form the foundation for an I-semantics of language, if we are willing to take the next step. And nobody ever suffered from learning and knowing a new language, even a technical one, or learning to structure an analytic problem in a different way. Also, there are many people who have been working on I-semantics informally for years whose work and insight is now ripe for inclusion in the new unified research program.

⁷ When it appears it appears first without the helping *be* verb, and it seems to occur first in telic verbs and then is gradually extended to verbs without salient endpoints. It is never apparently overextended to stative verbs.

2.4 | Reconstruction

In the reconstruction stage, the grieving semanticist turns her mind to productive engagement with the new. In this section, I wish to engage more closely with one particular aspect of Pietroski's positive proposal, specifically where it applies to the semantic combination of verbal events and their arguments.

Pietroski's view of the way participants are semantically integrated with the event is that the verb itself denotes a monadic property and that there are special dyadic properties, which look a little like neo-Davidsonian thematic roles, which allow participants to "clip on" to the event property in systematic ways. This constrained and judicious use of dyadic properties for precisely this kind of case gives a lot of extra combinatoric power to the system, without unleashing the full power of the lambda calculus.

My argument in this section will be that once the fine-grained details of the linguistic semantic generalizations about events are understood, and the crosslinguistic universals are seen, then an alternative view becomes more compelling.

In order to argue this, I need to digress to briefly lay out and summarize what we know about sub-evental structure in natural languages and how it is lexicalized by verbal symbols. I too believe that we need to build a minimalistic semantics based on the actual cognitive combinatoric relationships that languages seem to employ. I agree that monadic property predication is the heart of the system of verbal combinatorics, but instead of thematic role dyadic properties, I would like to argue that *causation between subevents* is all we need in addition to Aristotelian predication.

The first important distinction is the State versus Dynamic Event distinction, and in most languages, this comes with a robust list of criterial behaviors and diagnostics. Concerning states, the participants that are involved conform to a rather simple semantic notion—the Figure (in the terms of Talmy, 2000) holds the property described by the stative predicate. Or as Talmy himself puts it:

(13) The Figure-Ground Asymmetry:

The Figure is a moving or conceptually movable entity whose path, site, or orientation is conceived of as a variable, the particular value of which is the relevant issue.

The Ground is the reference entity, one that has stationary setting relative to a reference frame, with respect to which the Figure's path, site or orientation is characterized. (Talmy, 2000)

The Ground is the element that forms part of the predication and when it is itself a nominal, functions not as the thing whose properties are being asserted, but the thing acting as a reference object for another's properties. This Figure-Ground distinction is an important cognitively central core to the asymmetrical nature of all simple predications. We can extend the definition of Figure/Ground from the purely spatial domain to encompass stative properties more generally: the Figure of a property predication is the entity whose degree of possession of a particular property is at issue; the Ground is the reference property, or property scale which the Figure is predicated to "hold" to some degree.⁸ Interestingly, this predicational asymmetry corresponds to a syntactic one, with adpositional elements overwhelmingly, and possibly universally selecting for Grounds as complements (see Svenonius, 2010 for discussion), with the Figure as the notional "subject" of the relation (Svenonius, 2010; Talmy, 2000).

⁸ Kratzer (1996) uses the term "holder" for the introduction of and use of this general role label. With the exception of the "experiencer" role, fine grained differences in thematic role are not usually proposed for the stative Subject.

Taking the core difference between dynamic eventualities and states as our starting point, the minimal dynamic eventuality is one that characterizes most purely a simple event of dynamicity/change/process as opposed to the description of a static state of affairs. Let us simply represent these as two primitively different kinds of eventuality and notate them as e_d and e_s for convenience in what follows. Complex eventuality types can be built from these building blocks, but empirically, such causal complexity seems to be restricted in natural language. In fact, within the expressive remit of a single lexical verb, the typology is exactly what one would predict to exist if one could only do causal composition, and where the ingredients involve no more than one core dynamic eventuality, plus optional causationally upstream and downstream states (14).

- (14) DYNAMIC EVENT: e_{dyn}
 CAUSED DYNAMIC EVENT: $e_{cause} \rightarrow e_{dyn}$
 DYNAMIC EVENT WITH RESULT: $e_{dyn} \rightarrow e_{result}$
 CAUSED DYNAMIC EVENT WITH RESULT: $e_{cause} \rightarrow (e_{dyn} \rightarrow e_{result})$

Examples of these types in English are given in (15) below, for concreteness.

- (15) Lexical Verbs in English Across Event Types
- | | [+durative] | [-durative] |
|---|--------------------------|--------------------------|
| <i>Dynamic event</i> | Rise | Disappear |
| <i>Caused dynamic event</i> | Raise | Hit |
| <i>Caused dynamic event with result</i> | Destroy | Win |
| <i>Dynamic event with result</i> | empty _{intrans} | break _{intrans} |

Once again, there is a robust and exceptionless *syntactic* correlate to causal complexity, expressed in (16).

- (16) Event Complexity Generalizations:


The causing event, when it can be seen to be explicitly added, always adds morphology or participants that are hierarchically *above* the core dynamic event; result events are always added *below* the core dynamic event. Thus, the Cause head when it is invoked in the syntax is always on top of the main V (Folli & Harley, 2006; Pytkänen, 1999), and the result projection when added is always downstream of the main V (Hoekstra, 1988).

The next point is crucial, we also know that there is a robust argument structure hierarchy found across languages. The Undergoer and or Path argument is always hierarchically lower than the Causer/Initiator argument when both are expressed as direct arguments. As far as we can tell, both the event complexity generalization concerning causativizing morphology and the argument structure generalization seem to be exceptionless. But now, it only remains to notice that the *argument structure and aktionsart generalizations are in lockstep!*

- The Cause event is associated, when it exists, with an *external* argument.
- The Result event either introduces a new *internal* argument or is constrained to modify it (Levin & Rappaport Hovav, 1995).


- Arguments of the Change property can be either *internal* or *external* depending on what other arguments are present.
- Both argument structure and event structure of monoclausal verbal projections are constrained to express no more than one independent path of change.

The generalization seems clear: hierarchical representation of argument positions correlates with the hierarchical position of the subevent they are related to. To capture the pattern, and given the independent cognitive plausibility of causal complexity corresponding to embedding, all that is necessary is a notion of *Generalized property predication* as the universal semantic glue that relates verbal meaning to nominal meaning. And here of course, we do not mean predication in the sense of anything that you could build with the lambda calculus, but rather basic Aristotelian property predication. We must assume a general cognitive primitive of property ascription, for cognitively discriminable entities. This, we should assume, is syntactically represented at the “same” hierarchical level as the property description itself; traditionally we place it syntactically in the “specifier” position of the head.



We need to generalize just one step further to the dynamic version of this idea— if a property can be predicated of an individual, then a *changing* property can also be predicated of an individual. This creates the Undergoer relation, and it has the same specifier complement structure that the stative predication above has, with the only difference that the predicational head here is dynamic. The complement of a predicate of Change is Path.

(17) Dynamic Property Predication (± continuous)



We can now take these simple minimal predicational types and use the relations among subevents to build hierarchical subevental structures, where subevental embedding corresponds as a matter of general principle to the *cause/leads-to* relation. Each subevental structure is related potentially to a participant NP via property predication (either static or dynamic).

Syntactic details aside, my dissatisfaction with the proposal presented in Pietroski is not a dissatisfaction with the general agenda, but with the decision to implement the event domain using primitive contentful thematic predicates. In my own “minimal” version, I do not assume a classification of relation types independent of event structure relations and primitive property predication. I also do not have to assume a separate selectional or matching relation between the meaning of the verb and the thematic roles that is compatible with (but see Collins, 2020 for a discussion of this aspect of Pietroski's proposal). In my view, the verbal symbol offers a web of conceptual information corresponding to the identification of a (potentially complex) eventive structure, but not thematic relations per se.

The maximally expanded subevental structure for caused changes leading to a result, consists of a stative predication embedding a dynamic one, and the dynamic one in turn embedding a stative one. This structure and its proper subsets are intended to exhaust the event types built by the event building phase of the grammar. Thematic roles do not need to be listed separately, nor do their properties need to be memorized or known in advance. Interpreting phrasal embedding as causation will ensure the relative prominence of the different argument positions, and the minimal relationships of property-holding (both static and dynamic) will derive specific and different entailments for the different positions.

The labels on the trees above should therefore be seen not as labels in a template, they are there for ease of readability. The functional sequence here is actually quite spare, once the effects of causal embedding and predication are factored out. I think that this version of the event calculus is just a “minimal” as the Pietroski one. It has roughly the same amount of deviation from monadicity (causation between subevents rather than thematic relations), but it puts the basic notion of Aristotelian predication into a more prominent position in the combination of event with its arguments.

2.5 | Acceptance

I close this commentary with a reiteration of what I see as the central value of Pietroski's new book. With *Conjoining meanings*, he has made it impossible for anybody to think that there is no legitimate formal alternative to the traditional formal semantic toolbox. On the contrary, there are many reasons to think that it should be abandoned. Importantly, the philosophically deconstructive work done in the book is also paired with numerous interesting and original specific proposals for filling out the conjunctive agenda. I think the principle of minimalist internalist semantics is even more important than the individual proposals that Pietroski makes, although I do think it is important that he makes them. If this book is to have the impact that it deserves to have, it needs to be accepted by other working semanticists, and work on solutions within the new paradigm needs to become a joint enterprise of the many, instead of the lonely thankless business of the few.

ACKNOWLEDGMENTS

I thank John Collins and two anonymous reviewers for comments on an earlier draft of this article.

ORCID

Gillian Ramchand  <https://orcid.org/0000-0001-8019-0630>

REFERENCES

- Bennett, M. and B. Partee. (1978). *Toward the Logic of Tense and Aspect in English*. Reproduced and distributed (with added postscript) by the Indiana University Linguistics Club.
- Brown, R. (1973). *A first language: The early stages*. London: George Allen & Unwin.
- Collins, J. (2020). Conjoining meanings without loosing our heads, *Mind & Language*, 35.
- Condoravdi, C. (2002). Temporal interpretation of modals. Modals for the present and for the past. In D. Beaver, S. Kaufmann, B. Clark & L. Casillas (Eds.), *The construction of meaning* (pp. 59–87). Stanford, CA: CSLI Publications.
- Dowty, D. R. (1979). *Word meaning and Montague grammar: The semantics of verbs and times in generative semantics and in Montague's PTQ*. Dordrecht: Reidel.
- Folli, R. & Harley, H. (2006). On the licensing of causatives of directed motion: Waltzing Matilda all over. *Studia Linguistica*, 60(2), 1–35.
- Hallman, P. (2009). Proportions in time: Interactions of quantification and aspect. *Natural Language Semantics*, 17, 29–61.
- Heim, I. & Kratzer, K. (1998). *Semantics in Generative Grammar*. Oxford, England: Blackwell Publishers.
- Hoekstra, T. (1988). Small clause results. *Lingua*, 74(2–3), 101–139.
- Kratzer, A. (1977). What “must” and “can” must and can mean. *Linguistics and Philosophy*, 1(1), 337–355.
- Kratzer, A. (1996). Severing the external argument from the verb. In J. Rooryck & L. Zaring (Eds.), *Phrase structure and the lexicon* (pp. 109–137). Dordrecht: Kluwer.
- Landman, F. (1992). The progressive. *Natural Language Semantics*, 1(1), 1–32.
- Levin, B. & Rappaport Hovav, M. (1995). *Unaccusativity: At the syntax-lexical semantics Interface*. Cambridge, MA: MIT Press.
- Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.
- Marr, D. (1982). *Vision. A computational investigation into the human representation and processing of visual information*. San Francisco, CA: W. H. Freeman and Company.
- Owens, R. E. (2001). *Language development: An introduction* (5th ed.). Boston, MA: Allyn and Bacon.
- Parsons, T. (1990). *Events in the semantics of English: A study in subatomic semantics*. Cambridge, MA: MIT Press.
- Portner, P. (1998). The progressive in modal semantics. *Language*, 74, 760–787.
- Pylkkänen, L. (1999). Causation and external arguments. In L. Pylkkänen, A. van Hout & H. Harley (Eds.), *Papers from the UPenn/MIT roundtable on the lexicon* (pp. 161–183). Cambridge, MA: MITWPL 35.
- Ramchand, G. (2018). *Situations and syntactic structures: Rethinking auxiliaries and order in English. Linguistic inquiry monograph 77*. Cambridge, MA: MIT Press.
- Svenonius, P. (2010). Spatial P in English. In G. Cinque & L. Rizzi (Eds.), *The cartography of syntactic structures. Mapping spatial PPs* (Vol. 6). Oxford, England: Oxford University Press.
- Talmy, L. (2000). *Toward a cognitive semantics*. Cambridge, MA: MIT Press.

How to cite this article: Ramchand G. Truth is dead; long live the truth. Commentary on *Conjoining Meanings* by Paul Pietroski. *Mind & Language*. 2020;35:251–265. <https://doi.org/10.1111/mila.12269>