

# Recursion and Truth

Wolfram Hinzen

**Abstract** The paper describes reasons for departing from a traditional ‘T-model’ architecture of grammar and moving towards the ‘no interface’ view, on which grammar, inherently, organizes certain kinds of semantic information which are hypothesized to not exist in the absence of grammatical organization. In this sense, these are forms of *grammatical meaning*, and the claim is that there is no such meaning in some pre-linguistic ‘semantic component’, located on the other side of some ‘interface’. The paper particularly articulates this view with regards to reference and truth as inherent aspects of grammatical semantics, and argues that apparent constraints on recursion fall out from the way in which grammar organizes the forms of intent/sional reference that we universally find in human language.

**Keywords** Recursion • Truth • Reference • Grammatical semantics • X-within-X • Intensionality • Interfaces

## 1 The Origins of Recursion

On a common understanding, human language is an instance of a large range of systems that exhibit a discrete infinity of hierarchically organized objects. If we model each of these objects mathematically as a binary (two-membered) set and the operation of creating such a set is recursive, the sets will contain other such sets, until we hit lexical items, at the bottom of this system. These are its atoms: elements exhibiting no hierarchical syntactic complexity at all. Regarding the origins of atoms in hominin mental evolution – some mechanism for the encapsulation of

---

W. Hinzen (✉)

Catalan Institute for Advanced Studies and Research (ICREA), Barcelona, Spain

Department of Linguistics, Universitat de Barcelona, Barcelona, Spain

e-mail: [hinzen@ub.edu](mailto:hinzen@ub.edu)

complex conceptual structures in a non-structured primitive element – little appears to be known (see Uriagereka 2008 for some discussion). Regarding the origin of recursion, nothing much appears to be known either. Recent research has rather raised it to the level of the one significant mystery that distinguishes human language and mentality from animal systems of thought and communication (Hauser et al. 2002). Often the thought has been that the basic operation defined to generate the above structure (the operation Merge of recent minimalist theorizing: see e.g. Chomsky 2008a), comes ‘for free’, in the sense that it is minimally required for *any* system with the basic features of language. Put differently, it is ‘virtually conceptually necessary’. But this is not an explanation, any more than Kant’s claim that, in order for scientific knowledge to be possible, our minds need to have certain ‘categories’, provides an explanation for how or why our mind has that organization.<sup>1</sup>

Modelled as above, moreover, recursion tells us no more about human language than about any other hierarchically organized, discretely infinite system – from organic growth (Prusinkiewicz and Lindenmeyer 1996) to music (Katz and Pesetsky 2009) to DNA (Searls 2002), if not to a large range of other human and non-human systems of planning, foraging, and navigation. Whenever a system has discrete infinity, Merge can be employed to apply to whatever its atoms are and to form sets from them that contain other sets. And unless we block the ‘internal’ application of Merge (its application to sets and lexical items contained in other sets), the system will have ‘transformations’ as well. So, linguistic specificity will not come from Merge. It is also not clear why much specificity would come from the use of some of its atoms as ‘labels’ of the resulting sets (Hornstein 2009), if indeed labels are required for the workings of the linguistic system at all (see Chomsky 2008b).<sup>2</sup>

Nowadays, the standard answer to the question of what leads from Merge to the kind of recursions specific to human language is interface conditions, leaving Merge itself to apply completely freely (see e.g. Boeckx, this volume). That is, how recursion is constrained follows from constraints externally imposed on Merge by language-external systems, particularly the conceptual-intentional (C-I) system. This picture is broadly in line with (though not tied to) a traditional view on which syntax is ‘autonomous’, ‘formal’, and ‘modular’, whereas semantics is external to it and independently motivated. An alternative is that there is no distinctive syntax-module at all, any more than a distinctive semantics module – and hence

---

<sup>1</sup>Chomsky (2010) also argues that the question of precursors for Merge is mute – hence it evolved all at once: ‘There are many proposals involving precursors with Merge bounded: an operation to form two-word expressions from single words to reduce memory load for the lexicon, then another operation to form three-word expressions, etc. There is no evidence for this, and no obvious rationale either, since it is still necessary to assume that at some point unbounded Merge appears. Hence the assumption of earlier stages seems superfluous.’ While this may well be plausible, it doesn’t make the problem easier.

<sup>2</sup>Thus, to construct a purely hypothetical example, one might propose that the Merging of a Path with a Goal in insect navigation is still a Path, namely a bounded one: the Path ‘projects’ in this sense and becomes the ‘label’.

there is no C-I interface between them either, though there will be relevant systems of discourse processing that process the outputs of the language faculty further (Hinzen 2006; an idea that Chomsky 2008a, describes as a ‘radical’ one).<sup>3</sup> There is only one system and it generates all the distinctions we need. In this case, constraints on recursion need to follow from the workings of this one system itself: they can’t be externally imposed.

The present paper pursues this ‘radical’ model, building on Arsenijevic and Hinzen (2013). The latter paper observes apparently universal and somewhat surprising *restrictions* on recursion in human language. Specifically, (i) Universal Grammar appears to ban the occurrence of any syntactic category immediately in itself, i.e. the configuration  $X[X]$ , where  $X$  is any category: typically, the two occurrences of  $X$  are kept distinct through an interleaving category  $Y \neq X$ .<sup>4</sup> Moreover, (ii), where a category  $X$  embeds in itself indirectly (via interleaving of other categories), the embedded XP behaves differently from the embedding one, hence strictly there is no recursion (in the linguistic sense of *self-embedding*) in this case either.<sup>5</sup> I return to this view here, tracing its consequences for semantics and compositionality more specifically. In the Arsenijevic and Hinzen model, the unit of recursion is the cycle: it is the primitive on which recursion in human language is based. As viewed in the phase-based model of syntax (Chomsky 2008a, b), the cycle is inherently also a unit of semantic evaluation. Syntax thus revolves around semantic distinctions that correspond to different cycles. In this sense it is organized around semantic principles, and recursion is a consequence of the way in which syntactic derivations subserve semantic distinctions (specifically, semantic evaluations at the syntax-discourse interface).<sup>6</sup> This has consequences, I aim to show here, for the naturalization of truth, the fundamental primitive of semantic theory.

There is no dispute about the fact that speakers evaluate some linguistic expressions as true or false – a contingent aspect of human semantics which a recursive, systematic, and compositional form of communication or thought does not need to exhibit (Carstairs-McCarthy 1999). It therefore needs to be explained: we need to understand why the language system exhibits this specific form of organization.

---

<sup>3</sup>This architecture is in stark contrast to a system such as Jackendoff’s (2002), say, where the autonomy of syntax remains axiomatic. See e.g. p. 124: ‘meaning has an inventory of basic units and of means to combine them that is as distinct from syntax as syntax is from phonology . . . we ought to attempt to evacuate all semantic content from syntactic structure, just as we removed phonological content’. The present view is not in conflict with the old and valid insight that, as Lasnik (2000: 186) puts it, ‘there is no rule of *red-extraction*’: a syntactic rule constrained to apply only when a lexical item with the content RED is inserted.

<sup>4</sup>See also Roeper and Hollebrandse (2008) on recursion being typically ‘indirect’.

<sup>5</sup>None of this is in conflict with language being recursive in the sense of computable (though I don’t regard the computability of language as an empirically established fact either).

<sup>6</sup>This is not in contradiction with the above view that there is no C-I interface as on the standard minimalist view C-I systems are richly (in particular propositionally) structured, and they impose systematic forms of organization on the syntax. On the view developed here, there is a discourse representation, but it is ‘passive’ and consists of a set of discourse referents (which are then updated by syntactic structures) only.

The core idea explored in this paper is that the answer to this question is actually linked to the question about the origin and nature of recursion. The naturalization of recursion and of truth should proceed together.

## 2 A Correlation Between Truth and Recursion

I begin by observing some correlations between truth (as a specific form of semantic evaluation at particular moments of a derivation) and recursion. As I argue in this and the next section, there is recursion (in the sense of hypotactic embedding: coordination and recursion in adjunction will not be discussed here) only and as long as expressions are not evaluated for truth. Semantic evaluation for truth constrains syntactic recursions both so as to generate only objects of a certain kind and so as not to let these objects, once they are constructed, recur in themselves.

Note, then, that the paradigmatically recursive structures in language are the truth-evaluable ones, but that the range of truth-evaluable structures in language is in fact highly restricted.<sup>7</sup> To start with, it is clear and interesting that *words*, even though, crucially, they can be hierarchically, thematically, and compositionally structured, are not evaluable for truth. As di Sciullo (2005) discusses, words can encode some thematic (i.e., event) structure, as when incorporating an agent or theme (1a, b), even if not both (1c):

- (1a) tut-or
- (1b) tut-ee
- (1c) \*or-tut-ee

Compound words can resemble sentences and phrases in a similar way, as when ‘doll house’ denotes a house for a doll, a ‘doll house shop’ a shop whose purpose is the selling of houses for dolls, and a ‘doll house shop seller’ a person in the habit or profession of selling doll house shops. Words also can encode an ontological object-event distinction, as in (2), where the same lexical root,  $\sqrt{\text{GROW}}$ , is morpho-syntactically specified so as to denote a temporally extended activity located as simultaneous with the point of speech (2a), or an objectified process that cannot incorporate Agentivity (2b) (cf. *\*John’s growth of tomatoes*), or an objectified process that can (2c) (cf. *John’s growing of tomatoes*):

- (2a)  $\sqrt{\text{GROW}}$ -s
- (2b)  $\sqrt{\text{GROW}}$ -th
- (2c)  $\sqrt{\text{GROW}}$ -ing

---

<sup>7</sup>In ‘type-driven’ approaches to semantic compositionality this is often not observed. Thus, it is common to find that VPs that contain a subject position are said to be ‘semantically of type t’ (e.g. Heim and Kratzer 1998: 217). While this is justified to the extent that it makes the recursions run smoothly, VPs don’t e.g. contain temporal information, which any evaluation for truth in natural language requires.

Finally, as di Sciullo also notes, words have, again like sentences, a systematically compositional semantic structure, as the difference between (3a) and (3b) suggests:

- (3a) tavolo-in-etto 'funny little table'
- (3b) tavolo-letto-ino 'little funny table'

Against all of this semantic potential of words stand their intrinsic semantic limitations. Thus no English word can refer to truth-values: 'John' e.g., cannot denote one:

- (4) \*John is true

Relatedly, a word, no matter how complex and even if configuring an event or thematic structure, cannot anchor that event in time: (5a) is incapable of expressing what (5b) does, in which the (compound) word is replaced by a phrase (di Sciullo 2005: 51):

- (5a) John is a bank robber
- (5b) John is robbing a bank

Nor is (5c) possible, showing that however complex the semantic structure of a word becomes, it retains a curious generic character:

- (5c) \*At six, Mary is a writer

Relatedly, while words can expand productively, their expandability is inherently bounded in a way that prevents us from talking about a concrete instance of what abstract kind they denote (example from di Sciullo 2005: 36):

- (6) constitute-ion, constitute-ion-al; constitute-ion-al-ity;  
anti-constitute-ion-al-ity; \*th-anti-constitute-ion-al-ity

Thus, words cannot encode predications, and for the same reason, no compound word can say, of some person, that he is in the habit of selling doll house shops.

These limitations even show in the use of proper names. Thus, the name 'John' is not only generic in the sense that millions of people have it, but it is also generic in the sense that when used to denote one particular individual familiar to the interlocutors, it cannot denote John at this or that moment, in this or that time of his life, in this or that state of his mind. That is 'John' denotes John generically, hence in a more abstract way, disregarding the vagaries of his actions, mental state, time, and place. Similarly, Russell referred to simply as 'Russell' is not the early or the late Russell, Russell as a logician or as a lover, he is just Russell. Nor is a name, in and of itself, even capable of fixing what it is that it refers to. Thus, 'John' in 'John stands here' might denote his car; 'Johns' will denote persons having this name; and 'John was all over the floor' might denote his flesh, after a massacre (Hinzen 2007: ch. 4). Generally speaking, names need a phrase or sentence before their reference is fixed and fully specific. Thus, (7), where the name 'Nixon' occurs in a compound,

- (7a) John is a [Nixon-admirer] in every respect except that he does not admire Nixon
- (7b) ?John admires Nixon in every respect except that he does not admire Nixon

paradoxically shows that a ‘Nixon-admirer’ is not necessarily someone who admires Nixon (di Sciullo 2005: 52):

Neither can names that are parts of compounds be targeted for extraction (8), unlike in the case of phrases (9):

- (8a) John is a [Nixon-admirer]
- (8b) \*who is John an [t admirer]?
- (9a) John [admires Nixon]
- (9b) who does John [admire t]?

The right conclusion therefore appears to be that far from being the paradigms of referentiality, words as well as even names only attain referential specificity and definiteness, in short full referential potential, *after* entering phrases and sentences, i.e. syntax. Syntax mediates semantic reference. At the level of a lexical root such as  $\sqrt{\text{HOST}}$ , where the syntactic derivation has not yet begun to operate, it is not even clear yet whether we will use this root to refer to an object or an act (as in *the*  $\sqrt{\text{HOST}}$  versus  $\sqrt{\text{HOST-}ing}$  *some guests*); and although this particular semantic distinction is attainable at the level of compounds, we still cannot yet relate the abstract concepts of our minds as encoded in such compounds to specific moments in time and discourse, as we have seen, let alone put forward a proposition as true or false.

The widely assumed axiom in the philosophy of language and formal semantics, that semantic interpretation begins from the assignment of referents to *words* and the meaning of a sentence is built up ‘compositionally’ from there, thus needs to be qualified. Referential semantic interpretation is delayed at least until the phrase, if not the phrase, as we will argue later. As we have just seen, the compositional aspects of words and sentences are also logically *independent* of their referential aspects (cf. Uriagereka 2008): we could have the former in the absence of the latter (or at least the absence of referential specificity), and there could be a language where all expressions are semantically compositional (as words and compounds are), yet no specific object could ever be referred to with an expression of this language. In human language, reference only *gradually* arises as the syntactic derivation proceeds and relevant levels of categorial complexity are reached. This is true even for names, as we have noted, which go far beyond the ‘sticking of labels to objects’ and only in fact function as names when the syntax cooperates. Arguably, namehood presupposes full DP-syntax, explaining why naming is persistently absent in non-human animal calling and apparently more taxing in human cognitive processing than descriptions are (Semenza 2009).

As it turns out, in fact, the potential for an evaluation of truth arises only towards the very *end* of the syntactic process. Thus, no Noun Phrase, no matter how complex, can grammatically be said to be true or false, owing primarily to their lack of Tense, e.g. (10):

- (10a) Mary’s smile
- (10b) Caesar’s destruction/destroying of Syracuse

Neither need clausal structures exhibiting Tense have what it takes to be true or false:

- (11a) [IP to be an idiot]
- (11b) [CP for John to be an idiot]
- (11c) [CP if he is handsome]
- (11d) [CP that endorses the proposal]

Clauses used as root phrases need not qualify either:

- (12a) You idiot!
- (12b) Good cook, your mother!

Thus, (12a) is bound to an expressive rather than propositional meaning (Potts and Roeper 2007), and (12b) is evaluative as well, in a way that a proposition need not be. Neither of them can be hypotactically embedded in the way propositions can be, hence give rise to recursive structures. The same is true of small clauses, as (13) shows:

- (13a) \*[CP John considered [SC Joan see [SC her fall]]]
- (13b) \*[CP John thought [SC Joan happy [SC Mary sad]]]

In both cases, this is a consequence of their structural impoverishment. Again, Tenses are needed to generate recursions productively at the clausal level (14a), recursions that start running smoothly and productively only when these Tenses are finite ones in the scope of a subordinating conjunction:

- (14a) John [believed [Joan to consider [Mary to be happy]]]
- (14b) John [believed [that Joan thought [that Mary was happy]]]

In short, *both* truth and recursion come into their own once we have built structures that are, in some sense, ‘complete’ enough: specifically, we need a fully sequenced hierarchy of functional projections up to the lower regions of the C-field as described in Rizzi (1997), which is where we have a truth-evaluable syntactic object. Before we are there, productive embedding at the clausal level won’t happen. That in itself points to our conclusion that truth and recursion are inherently correlated in human language design, and that recursion is not ‘free’. What recurs in language, in paradigmatically recursive structures usually used to illustrate recursivity (e.g. *Sam believed John thought Bill knew Mary was pregnant*), are *truth-evaluable* structures. The other paradigmatic case of recursion (in hypotactic contexts) is recursion at the level of referentially evaluable objects, as in (15), a case to which I return:

- (15) the vase on the table in the room in the country house in France

Syntax therefore not only mediates reference. Truth and reference also mediate syntactic recursion: here again, the units of recursion are objects which are


semantically sufficiently complete, so as to be usable for purposes of reference. Rather than being free, recursion waits until that completeness is reached (or sufficiently approximated). Syntax is in this sense constructed around semantic principles of organization.

### 3 More Correlation

Let us now note that C by itself, e.g., is simply not recursive: C does not embed in C, directly, but only after the embedding has unrolled sufficiently, so that a whole sequence of projections intrinsic to the C-domain is in place (cf. Hollebrandse and Roeper 2008). Thus, e.g., (16) is ruled out, what we get instead is *cyclic* recursion, as in (17a) or (17b):

- (16) [C[C[C]]] ... ]  
 (17a) [C-T-V [C-T-V ... [C-T-V]]]  
 (17b) [C-T-V-D [C-T-V-D ... [C-T-V-D ... ]]]

To be specific now, what productively recurs appears to need to be at least as complex as IP plus perhaps the lower regions of CP as described in Rizzi (1997):




What comes above FinP here is elements governing the embedding of a propositional unit configured at the level of IP/Fin into the discourse. But recursion not only happens with units that are sufficiently complex, it also happens only as long as these units are not *too* complex, and embedding into the discourse has not proceeded too far. As soon as propositions have assertoric Force, in particular, they don't embed: embedded propositions are never asserted (although their truth values can be assumed in discourse, or be presupposed by the speaker, as in factives).<sup>8</sup>

<sup>8</sup>Note that in adjunct structures clauses with distinct assertoric forces *can* occur embedded: cf. *John said that the dog, which is really big, is small*. Here *that the dog is small* is assigned no truth value by the speaker, but *which is really big* is. Thanks to Tom Roeper, p.c., for the example.



If we terminologically distinguish between truth-evaluability and truth-evaluation (or assertion: the online assignment of a truth-value at a moment of speech), and associate the latter with ForceP, a constraint in human grammar design thus seems to be that the configuration in (19) is not possible:



That is, a truth-evaluated object can never recur in itself.<sup>9</sup> The empirical observation is therefore that upon reaching appropriate complexity levels which correlate with evaluability for truth, the generated syntactic object can recur; if it goes beyond this level of complexity, recursion stops: the object cannot recur. Not only Force is not recursive: neither are evidentiality markers that can be added at this point, or tag-questions in languages like English or discourse particles in a language like German with which the truth assigned to a given proposition can be negotiated with the hearer. The ‘purpose’ of a syntactic derivation, then, is to compute a truth-value. Once this is done, the recursion is halted. And compositionality is halted too: the truth-conditions are fixed by the time the FinP is computed; evidentiality markers, tags, etc., don’t *add* to the truth-theoretic content, on a common assumption. The same observation may also explain why the ‘it’s that’ construction, as Andrea Moro points out (p.c.) does not iterate<sup>10</sup>:

- (20) It’s that John is sad
- (21) \*It’s that it’s that John is sad

In sum, the moral is that recursion in human language is highly restricted, that these restrictions are of a semantic nature, and that the notion of truth has to be brought in to account for them. Linguistic recursions are not only bound to the

---

<sup>9</sup>One might note that, for probably related reasons, the referential Force specified at the left edge of the DP cannot recur either (cf. remarks in Higginbotham 1985: 560–3).  
<sup>10</sup>On Moro’s analysis, the structure of the latter sentence is this (cf. Moro 1997: 193):

[<sub>IP</sub> it<sub>i</sub> is [<sub>SC</sub> [<sub>CP</sub> that John is sad] t<sub>i</sub>]]

Here, *it* is base-generated in a Small Clause (SC) complement to the copula ‘is’, where it functions as a pro-predicative element that attaches as a predicate to the small-clausal subject, the clause ‘that John is sad’. When this element is raised (leaving the trace *t*), the result is interpreted as involving the assignment of a sentential predicate to the embedded proposition. Basically, this is the truth-predicate: (20) essentially means *It’s that it’s true that John is sad* (or *It’s the fact that John is sad*). It doesn’t mean anything else that it might logically speaking have meant, e.g. *It’s that it’s probable that John is sad* or *It’s that it’s highly certain that John is sad*. This suggests that truth, and no other form of semantic evaluation, is indeed what inherently (or by default) ‘closes off’ a given proposition. If truth evaluation is not recursive, or functions as a recursion-stopper, we predict the oddness or ungrammaticality of (21).

cycle, they also stop altogether when semantic complexity has reached the level of the truth-value. Syntactic derivations are mappings from lexical roots, where referentiality is maximally unspecified or minimal, to a fully projected clause, where it is maximal. A derivation is the computation of an extension from a given intension, which in turn is a pure concept not evaluated for reference.

There is no moment in the derivation on this model where syntax is ‘autonomous’, or language is arbitrarily ‘recursive’. Everything is bounded. It is only when we abstract from the inherent *uses* of specific syntactic configurations, that language appears to be generated by an unbounded recursive operation. Within the cycle, surely, everything is rigidly hierarchical and asymmetric: nothing ever recurs in it. We don’t get structures like (22) or (23), say:

(22) C-T-T-T... v-V

(23) C-T-C-T-C-T... v-V

The reason again is plainly semantic. The same is true within smaller cycles, like v-V itself. Where ASPECT phrases embed one another, for example, the relevant aspectual projections are of different kinds, and the same ones *don’t* seem to immediately recur (cf. Den Dikken 2009). Where VPs recur, as in serial verbs, a closer look reveals that the constructions in questions are either asymmetric in their supposedly directly recursive constituents (e.g. one being an argument, the other an adjunct), they involve different categories (as they certainly would on a ‘cartographic’ approach), or they are clearly templatic (Pawley 2006). vPs can systematically embed other vPs, as in complex causative constructions (cf. *cause to die*), but this process is not only highly limited, but the embedded vP is also never quite as complex or projected as the embedding one (semantically an event embeds a *state*, say, not another event of the same sort). Moreover, if an event-encoding vP contains sub-events, these are semantically necessarily *parts* of the larger events, not independent events: a complex vP will always denote a *single* event. NPs recur productively in themselves, but also only with other material interleaving (cf. *\*the mother the bride*). Even nominal root compounding, perhaps the next candidate for ‘direct’ recursion (‘X-in-X’), are plausibly analyzed as involving lexical access in between two Ns, if indeed they are not reduced structures with more material intervening. We also know that argument structure is highly restricted rather than unboundedly recursive; and that no sentence can have more than one subject (see Arsenijevic and Hinzen (2013) for more discussion of all of these examples). In short, the ban on recursion in the above sense above appears to be quite robust.

Nothing happening within the cycle is recursive or unboundedly productive: starting a phase from a nominal head, we know we can project up to, say, D, and then recursion stops: a new cycle has to be begun. Similarly, as we have seen, when starting with whatever is the initiating head of the C-phase, we can’t significantly project beyond C. So operations building the cycle are essentially templatic: they are predictable given the choice of the initial head and its inherent potential. If we let an unbounded operation generate the cycle, we would therefore have to restrict it. If we begin from the restricted mental template, or *Gestalt*, by contrast, there is no unbounded operation to be restricted. If the cycle was recursive, or there would

be an infinite sequence of projections contrary to the finite one that Rizzi (1997) or Cinque (1999) depicted, there wouldn't in fact *be* the specific recursions that we find in human languages. Instead, like in the natural numbers, we would be looking at an infinite sequence of uniform objects, all of the same nature, with no cycles of the sort we find in language ever forming (or, perhaps, the only cycle there would be the single step from one number to the next) (cf. Hinzen 2009). The source of recursion in human language, therefore, is in fact the inherent *limitations* of language, or its *boundedness* – which is what the cycle reflects.

Evaluation for truth is thus woven into the dynamics of structure-building in language. A consequence of this is that truth, inherently, is a *structural* notion, which cannot be explained by anything like general-purpose 'relations to the world' or notions of 'correspondence', which seem non-explanatory and in fact to trivialize the problem. The problem of the origin of truth is not solved by positing a 'truth-relation' to the world in which certain abstract entities such as 'propositions' stand. Rather, syntax – and specifically the cycle – is an essential part of the answer to the question of the origin of truth (though certainly not a complete one: this is *not* being claimed here). An alternative answer to this question has been to reject it. Thus, one might maintain that truth is a paradigmatically externalist notion that falls strictly outside the domain of 'internalist' inquiry into the mind/brain as generative grammar has conceived it (cf. Chomsky 2000). If truth is primarily a relation to the world, this might indeed follow. But truth may not be a relation, and arguably, Tarski (1933) and more recent 'deflationist' theorizing in the metaphysics of truth has shown us how it might not be (e.g. Field 1994). Even if it is a relation, however, our question was not the nature of any particular truth, but the origin of truth. In that respect my argument has been that, quite clearly, evaluation for truth is woven into the fabric of syntax and it is an aspect of how grammar is recursively organized.<sup>11</sup>

## 4 Compositionality in Modern Semantics

The principle of compositionality is written into the modern conception of semantics. As discussed in Szabo (2009), a sensibly disambiguated version of the principle is that when we have (i) fixed the meanings of the syntactic constituents of an expression, and (ii) fixed its syntactic structure, its overall meaning is fixed as well: it won't depend on anything else. The first formalization of this idea appears to be due to Tarski (1933) (cf. Hodges 2001). Defining syntax precisely and without any reference to meaning, Tarski showed, for each formula  $\varphi$  of a formal language  $L$ ,

---

<sup>11</sup>A reference-theoretic approach to semantics, I have argued earlier (Hinzen 2007), confuses the problem with the answer to the problem: reference in an explanatory semantic theory is the problem, not the answer (and language is part of the answer). How we refer to the world in the way we do is exactly what we want to know. This problem is not answered by assuming that there is a 'reference-relation' between language and world; nor do causal relations seem to turn the trick, in the human case (cf. Hauser 2000).

how to define, by recursion on the syntactic complexity of  $\varphi$ , the class  $\mu(\varphi)$  of those assignments (of objects to the free variables of  $\varphi$ ) which satisfy  $\varphi$ . We see here precisely what we questioned above for the case of human language (with which Tarski was not directly concerned): recursion is defined without regard to semantic considerations. The question now arises how this meaningless syntactic structure is mapped to a semantic structure. Compositionality was (and is) the answer. The meaning-function  $\mu(\varphi)$  is given by a function  $\chi(\varphi) = (\mu(\delta), \mu(\theta))$ , where  $\delta$  and  $\theta$  are immediate constituents of  $\varphi$  and  $\chi$  only depends on the syntactic operations generating  $\varphi$ . Note that every constituent is assigned an independent meaning, which never depends on what constituent it is embedded in.

As used in the Montagovian tradition and specifically by Montague himself, semantics in this sense is essentially neutral as to what exactly meanings are. Married with philosophical assumptions on the nature of meaning, on the other hand, the idea by and large has been that language is *externally controlled* somehow (it equates with denotation). The base case of meaning is thus where  $\varphi$  is simply a name and  $\mu(\varphi)$  is an object, its bearer. With this assumption, contradicting conclusions reached above, standard introductions to the philosophy of language set out (e.g., Lycan 2008), and Szabo (2009) formulates a broad consensus when he remarks that it is ‘very plausible that the meaning of proper names is nothing but their referent (for the point of having proper names in a language does not seem to go beyond the labelling of objects).’

So much was naming (thus understood) the base case, that all *other* denotations that a compositional semantics will have to assign to syntactic constituents have been defined *in terms* of the denotations of names, i.e. individual objects: e.g., as sets of such objects (in the case of adjectives and common nouns), or as sets of such sets (as in the case of generalized quantifiers). A predicate denotation thus is a mapping from an object to an object (type  $\langle e, t \rangle$ , or a set of objects characterized by that function). Basically, then, every constituent is assigned a referent, and where these are sets they are construed as unary functions applying to objects in their domain of definition, until a single individual referent results, a truth-value. In a sense, then, the ontology of semantics never goes beyond the base case, the ontology of individuals. To put this differently, the ontology of semantics, as standardly conceived, is never *relational*: the only thing there is is individuals and set-theoretic constructions from them.<sup>12</sup> It is as if meaning could never arise from anything other than some sort of relation between a symbol and a thing. In short, although Tarski did formal work and that work is neutral on the nature of meaning, in actual practice it has been linked with a conception of meaning as reference, and every syntactic constituent is evaluated for reference (though some of these are functions/sets and most ‘referents’ are theory-internal constructs). The referents of a complex constituent are then recursively evaluated in terms of the referents of their immediate parts.

This conception, I will now argue, leaves a venerable Ancient problem unsolved. Suppose we wish to understand why (24) means what it does:

---

<sup>12</sup>This is a *nominalist* ontology (reflecting Occamist premises), as described in Leiss (2009).

## (24) Theaetetus sits.

The standard referential conception in semantics begins my mapping ‘Theaetetus’ to Theaetetus – so far so good.<sup>13</sup> And then ‘is wise’ is mapped to another referent, let us call it a ‘property’. So we have two referents: an object and a property. But the meaning, a proposition, is neither of these two. And it is not of the same kind as either of them. So it is something third. But then we have three objects or referents, and we are not closer to what makes for the unity of a proposition. So let us go back to the two first objects. What provides for their *link*? Plato posited a relation of ‘instantiation’ between them: the object ‘partakes’ in the property, which somehow embodies generality (though God only knows how). But this relation is again something third. We would have again three entities, and would again need something fourth that relates them, creating a regress leaving us with the same problem as before.

Plato’s proposal was critically reviewed for two millennia. It kept intriguing the greatest of minds in all the centuries that passed. The problem has remained: what makes for the transition between a name and what is not a name and neither a sequence of names, namely a sentence? Somehow, *saying* something does not appear to be the same thing as *referring* to something. Semantics is *more* than reference, and is not resolved by positing a new external object (‘properties’, ‘propositions’, etc.) whenever we need one for some non-name like constituent to refer to: ontology doesn’t save semantic theory.

How could an approach solve this problem in which anything that is relational is ultimately reduced to something that is not, namely names and set-theoretic constructions from them, as we have seen above? Set-theory is not the answer, even if we believed in the existence of sets, and could define the element-of relation consistently, and believed, unlike Pietroski (2003), that predicate extensions *are* sets. If compositionality is assumed, the problem aggravates: each constituent then independently refers, and before referents are composed, a sentence meaning is simply a set of referents. How does that set become a proposition? Looking at any referent assigned to a predicative expression won’t answer that question: all it tells us is that a certain object (say, Theaetetus) is mapped to another object (say, truth). But all this says is that ‘sits’ does what we know predicates do. How do predicates do that?<sup>14</sup> Again, positing ‘truth’ as another object in one’s ontology for sentences to refer to does precisely not explain what makes for the unity of a proposition, of

---

<sup>13</sup>Notwithstanding the fact no one quite knows what kind of object Theaetetus might be. Thus he is clearly dead, and dead objects are not somehow available still as objects of reference, like things stored in a warehouse. So is this referent a physical object? An object that once was a physical object, but is individuated irrespective of its physical instantiation? A physical object, when it still existed, at a certain mature time of its life? A person that as such is not a physical object?

<sup>14</sup>See Davidson (2005: 157–8) for discussion. The predicate ‘being a member of the seated things’ does not explain the original role of the predicate ‘sits’: it is a new predicate, for which the same problem arises. See also Higginbotham (2009: 149, for an illuminating discussion). The naïve idea that the extension of the common noun ‘dog’ is the set of individual dogs also ignores the fact that, in English, it is the word ‘dogs’ that refers to such a set, not the word ‘dog’. Given compositionality

which, as Davidson notes, the truth value is the sign: ‘only whole sentences have a truth value’ (Davidson 2005: 87).

Davidson’s brilliant late (2005) book is an admirable attempt to trace the history of this problem, which, in many ways, can be seen *as* the problem of the origin of truth. Davidson reads that history as a history of failed attempts, from Plato to Aristotle to Frege to Russell. Even Tarski, on Davidson’s account, did not solve it, despite the fact that he defined truth. Tarski’s achievement, on this reading, is to have told us how to set up a ‘T-theory’ for a relevant fragment of a given language, where T is a predicate applying to (structural descriptions of) sentences. T is defined such that, whenever a sentence *S* such as ‘snow is white’ is in its extension, snow is indeed white (and vice versa). More precisely, the theory is constrained so as to generate an infinity of sentences of this form (‘T-sentences’):

(25) *S* is T if and only if *p*.

Here for ‘*S*’ we may substitute a structural description of any sentence of the language for which we provide the theory (under certain restrictions), and for ‘*p*’ we substitute that very sentence (if our metalanguage or the language of the theory contains the language in question, as we may assume for present purposes). So, intuitively, T appears as a *truth*-predicate: for that is the predicate for which correct T-sentences would indeed intuitively hold (it is indeed so in English that, if ‘snow is white’ is true, snow is white). Moreover, and crucially, the definition of T tells us how the truth-conditions of sentences systematically depend on the meanings of their parts, hence it is *informative* in regards to how truth conditions are computed. It is our general concept of truth, however, the one against which we can test whether any given T-theory (or any given T-sentence) does indeed apply to a given language, which is *not* defined by Tarski. If, e.g., a T-sentence unlike (25) said that, in English, ‘snow is white’ is T iff grass is green, it would likely be a bad theory. Tarski presupposes our prior grasp of this concept. All Tarski told us, then, is how our knowledge of the conditions under which sentences are true systematically exploits our knowledge of the structure of these sentences. But all this is on the *basis* of our knowledge of truth and ability to apply this notion to utterances.

Tarski thus didn’t tell us how truth arises; and nor did he solve the problem of the unity of the proposition, or explain why sentences only have one truth value: this is what we will aim to improve on below. Put differently, it is not so that, when deprived of a notion of truth, Tarski’s recursive and compositional assignment of semantic values to parts of sentences would make us understand what truth is. On the contrary, as Davidson (2005: ch. 7) stresses, the assignment of a compositional structure to sentences is in the service of validating a given formal T-theory for a given language *L*, in the process of testing it against our truth-theoretic intuitions about *L*. The truth-conditions we assign to T are what we start from. They are our basic data. Once we have them, we work out how the truth values of wholes depend

---

we cannot specify the meaning of a morphologically less complex word in terms of that of a more complex one.

on their parts. Tarski's merit, on the other hand, is to have shown us that we cannot advance in the problem of truth on an ontological or reference-theoretic path: it does not help to posit an entity that a true sentence denotes; or to postulate truth values as objects for sentences to refer to; or to posit various referents for sub-sentential expressions as a basis for understanding the notion of truth. That would neither explain why we ever map a word to a referent, nor how a truth value arises from a set of referents: Tarski's semantics is not a reference-theoretic one, and I take that to speak in its favour. It is from here that what I have to say about the unity of the proposition departs. Tarski wrote long before modern linguistics got going, and philosophers started taking linguistic form seriously, as something that might explain us why it is mapped to meaning in the way it is. Here the previous section may provide a lead: recursions as visible in linguistic form are not arbitrary; and as I will now argue, they are not compositional in the above sense either.

## 5 Truth and Recursion, Again

Let us return to Davidson's talk of 'the truth-value of a sentence as a sign of the unity of a sentence: only whole sentences have a truth value' (Davidson 2005: 87). What this means is no more and no less than that there is no recurrence on truth: if a sentence is assigned a truth value, no part of the sentence, not even if it contains many other sentences, is assigned a truth-value too. Non-recursion and the problem of the unity of the proposition are thus closely related. The fact is that there is always only one truth value – no matter how many sub-sentences a sentence embeds, and no matter whether these sub-sentences could perfectly well be evaluated for truth when occurring as roots. Why should this be so?

From a traditional, modern semantic view as described above there is no reason at all for why this should hold, and nothing in semantics predicts it. Indeed, it is so puzzling that *intensionality*, which simply *is* the observed impossibility of assigning normal semantic values such as truth values as the denotations of embedded sentences, has been one of the most persistent and recalcitrant problems in semantics, seemingly invalidating the basic assumptions of semantic theory and the compositionality principle (see e.g. Higginbotham 2009: 143–4). From the viewpoint of compositionality, the meaning of a whole is a function of the meanings of the parts and their syntactic combination, but the meaning of the parts is *not* a function of the meaning of any whole that contains these parts. So a sentence that is an embedded constituent should contribute the meaning that it also has otherwise, in different structures, or when it occurs non-embedded.<sup>15</sup> If that meaning is a truth value, as it is on Frege's traditional view, we predict that the truth value of a whole is to be computed from (or determined by) the truth values of its parts. That, exactly,

---

<sup>15</sup>Frege's contention that in embedded contexts, sentences refer to senses rather than truth values, is in conflict with this principle.

is what we never find – no human language is like that. The truth value of a whole never depends on whether any of its parts are true or false. This is well known to be what takes children about 4 years to understand: before this time, in a situation where (26a) holds but Mom in fact bought oranges, they answer ‘oranges’ to the question in (26b) (de Villiers 2007):

- (26a) Mom said she bought apples.  
 (26b) What did Mom say she bought?

As the child approximates adult grammar, it learns that the truth value of a sentence like (26a) never depends in any systematic way on that of its embedded part. I interpret this as entailing three things: (i) only adult grammar analyzes complex sentences as cyclic, part-whole structures, in which (ii) the parts are never fully referential but always remain intensional; and thus (iii) referential force remains undivided and reserved to the structure as a whole. Put differently, the truth value bearing syntactic object is the ultimate non-recursive unit. It doesn’t multiply and it doesn’t occur in itself. Other things can occur in it, e.g. the denotations of embedded CPs, NPs, and *v*Ps, but semantically these are always weaker than a truth value: they can never determine one. Once the truth value is assigned at the root, the structure of these embedded parts gets fixated: compositionality and recursion stop. Whatever occurs embedded is thus to some degree intensionally interpreted: it is not referentially complete. This is obvious in the case of NPs: lacking even tense, they cannot serve to fully locate the denotation of the head noun of the NP in the world: the world is temporal as well as spatial, and no physically located object is without a temporal dimension. It is also clear in the case of *v*Ps: lacking Tense as well (though having Aspect), *v*Ps can only configure an event – but they cannot locate one in time and space, or in relation to the time and place of speech, for which the functional projections in INFL and higher up are needed.

None of this necessarily means that the principle of compositionality is wrong. As long as we factor into our account of it that whatever occurs embedded is destined to remain intensional, compositionality might be retained. However, assigning non-standard, intensional denotations would not only not explain intensionality. It would also not solve the problem of unity, and it would also not be how compositionality has standardly been understood (it was conceived in an extensional, reference-theoretic paradigm). The whole point of the compositionality principle is that embedded units have whatever meaning they have context-independently.<sup>16</sup> Only in this way can the meaning of a whole be composed from them. That embedded units should be intensional as a matter of principle makes no sense from this viewpoint. Why should hypotactic embedding and intensionality correlate?<sup>17</sup>

The answer seems to be that syntax matters: the meaning of a whole *cannot* be broken into parts: it is one single unit, which, like a *Gestalt*, has parts, but these

<sup>16</sup>See e.g. Fodor and Lepore (2002: 37): ‘the *compositionality* of complex concepts is of a piece with the *context-independence* of the contents of their constituents’ (emphasis in original).

<sup>17</sup>That, too, is an issue that Larson’s (2011) interesting approach fails to answer.



are, inherently, parts of the whole that contains them. This solution would solve the problem of the unity of the proposition, of which the truth value is the sign. It would predict why evaluation for truth is not recursive: it is no more recursive than a perceptual *Gestalt*, which also doesn't recursively occur within itself. The same solution would solve Frege's Puzzle about the intensionality of propositional attitude contexts. For the relevant reasoning here is this: In (27), the embedded sentence denotes a truth value; in (28), it also does, indeed in any possible world. But the intuitive truth values of (27) and (28) do not at all have to be the same:

- (27) Lois believes that Superman flies.
- (28) Lois believes that Clark Kent flies.

The answer now is: the assumption that embedded sentences denote truth values and that the truth values of their matrix sentences is computed from them is simply wrong: extensions are *never* composed in Universal Grammar. Only intensions are, until an extension is reached, in which case recursion stops.

Again we could say that *intensions* are composed, even if not extensions. So, even though 'Superman flies' might be more descriptive than referential, that description (or its semantic value) is still what the truth value of (27) is composed from. However, this would be to *concede* that the matrix CP and the embedded CP are not assigned the same semantic value: one is assigned a truth value, the other an intensional description. In a contemporary minimalist architecture, where functional projections govern the way syntactic objects function at the interface and are embedded in discourse, this would mean that they are not the same objects syntactically, and therefore again there wouldn't be recursion in cases like this either, if this entails the occurrence of the same syntactic object in itself, even if there would be compositionality. Moreover it would be puzzling why the two clauses behave so differently. We could also let the two clauses be syntactically the same and assign the same semantic object to them: a function from possible worlds to truth values in both cases. The function denoted by the embedded clause could then be composed with the denotation of the matrix verb, which would be defined so as to take the function as an argument. Then again however the two recursive clausal units would not be evaluated in the same fashion and we wouldn't know why this is the case: why should 'Superman flies' be a function from a possible world that is compatible with Lois' beliefs to a truth value, whereas the matrix clause is a function from the actual world to a truth value in that world? Reference would still shift when moving from the matrix to the embedded clause, and the assignment of the same semantic value to both would make it mysterious why this should happen.

One might try arguing that the reason that reference shifts is the specific nature of the matrix verb. Thus, could we solve the problem by assigning this verb a special semantic value, namely that of a relation between a person and a set of possible worlds compatible with what that person believes in the actual world, so that, e.g., (27) would be true just in case Superman flies in all worlds that are compatible with Lois' beliefs (see e.g., Heim and Kratzer 1998: 306)? Yet, that the matrix verb is indexed to the actual world whereas the embedded one is indexed to a merely possible one is exactly the fact we set out to explain: the solution now

considered restates that fact in formal terms. That a special principle, Intensional Functional Composition (Heim and Kratzer 1998: 308), needs to be invoked in these cases, seems to formalize the facts we needed to explain. Note, moreover, that the lexical semantics of the matrix belief-type verbs is as such entirely *consistent* with reference *not* being shifted in the above way. Semantically and compositionally, nothing speaks against the following interpretive rule:

- (29) (26a) is true if and only if Mom bought apples and Mom said so.

According to this rule, a speaker uttering (26a) would commit himself to two things: that the embedded clause is true and that the matrix one is true too. What he says would be true if he was right on both counts. The meaning of the word ‘says’ or ‘believes’ is as such compatible with such an interpretive rule. We shouldn’t blame the matrix verb for the intensionality effects if its semantics is entirely consistent with the absence of these effects. It seems more likely therefore that the problem cannot be solved in semantic terms and must lie in syntax: somehow it must be the case that as the syntax creates a part, it creates the intensionality effect as well. That would contradict compositionality as customarily understood, for the embedded unit would then, upon being embedded, precisely *not* any more mean what it would normally mean, and contribute that normal meaning to the compositional process. It would also contradict recursivity, insofar as an embedded part would never function as the whole that embeds it insofar as we take into account how syntactic constituents function at the interface. This would also make sense of the fact that the intensionality effects in hypotactic embedding are more general: arguably, it has nothing specifically to do with belief-type verbs and their lexical semantics. Thus, (30) clearly doesn’t mean the same as (31), not even if, in the context where (30) is uttered, the referent of ‘You’ is John Smith and the speaker knows that:

- (30) You are John Smith.  
 (31) John Smith is John Smith.

If the semantic value of ‘You’ and ‘John Smith’ was an extension, and it was the same object (as in the context above), the two sentences should have the same compositional meaning. But this is clearly not so. It makes for a difference whether we refer to a person grammatically as second-person or as third-person. Intensionality arises for grammatical reasons here, not for semantic ones. Similarly, suppose we refer to a particular object under the description in (32):

- (32) the vase on the table

In this case even though the reference of the embedded constituent ‘the table’ will need to be worked out in order to determine the reference of the whole expression, the embedded NP does not at all function on a par with the embedding NP. Certainly a speaker uttering (32) does not refer to two objects, as he would in (33), a paratactic context:

(33) this vase, this table

(32) is different from (33) in that ‘the table’ in (32) plays a quasi-thematic role (a location): it specifies a surface on which the vase is located. So it specifies the referent of ‘the vase’ further, and does not function as an independently referential unit. Indeed, if functioning as an independently referential unit, in a non-embedded context, it would precisely *not* denote what it denotes in (32): It would denote an individual, whole object, not a surface, as in (32), which is not the same as a table, but a particular aspect of it.<sup>18</sup> Moreover, for predicational purposes, the embedded NP in (32) behaves as if it wasn’t there: in *the vase on the table is beautiful*, say, nothing at all is said about the table. The contrast with (34) makes clear that we have a single referent in (32), but multiple independent referents in (34):

(34) This vase, this table \*is/are beautiful.

The reason must be syntactic. Hypotactic, recursive embedding as opposed to paratactic conjunction creates an intensionality effect, entirely independently of any attitudinal verbs or modal contexts.<sup>19</sup> Semantics is innocent; syntax is to blame. Syntax is the reality on which the principle of compositionality is silent.<sup>20</sup> In the remaining Section, I will turn to how syntax might create this effect.

## 6 Intensionality from Syntax

In its earliest formulation, in the mid-1960s, the ‘transformational cycle’ was a general principle on rule application: if there was a particular linear sequence of transformations acting on a single phrase marker, it would need to apply to the most deeply embedded sentential structure of the phrase marker first, before it could proceed to the next higher one. In the guise of the Strict Cycle Condition (SSC) of Chomsky (1973: 243), the cyclicity principle maintains that given an object as in (35) with two cyclic domains A and B, one embedded in the other, no rule can apply to such an object by solely targeting B:

(35) [A . . . . . [B . . . . . ]]

Freidin (1999), tracing the SSC’s history and writing from the perspective of early Minimalism, argues that it is a natural assumption that once a category has


<sup>18</sup>Or, as the case might also be, the floor under it, as in *the vase under the table*.

<sup>19</sup>Similarly, events embedded within others encoded in a complex vP need to be interpreted as parts of the larger events, not as independent events, as noted earlier, again an intensionality effect, given the intensionality of the notion of ‘part’.

<sup>20</sup>A principle that, both in Montague’s case and that of his contemporary followers never depended on any particular assumptions about syntax, or indeed whether there was such a thing in any real sense (e.g. Jacobson 2009).

become a constituent of another category (like B in (35)), it cannot project any more (Freidin 1999: 120). Its syntactic life is exhausted, as it were. Thus, in particular, no constituent within B can be targeted and moved to the left edge of B once (35) has been built. Things can move out of B only when B is ‘live’, and this is only as long as B is not embedded. One way of putting this is that the computation always only ‘sees’ a single cycle, the one in which it is currently working. By the time it applies to the object A, the cycle B has been completed and is not modifiable any more: Whatever plays a role in the further derivation must, by that time, have been moved out of B into its left edge so as still to be ‘visible’. All the derivation ‘sees’ when it goes to construct A, then, is the head of B and its left edge.

A cycle, then, is naturally conceived as a window of computational opportunity as well as memory, which is how it is conceived today (Chomsky 2008b): the derivation never looks back further than the left edge of the previous cycle, and when the cycle it is operating on is complete it is transferred out of the derivational workspace, with only its head and left periphery remaining, which then belong to the new cycle. In this model, the operation ‘Transfer’ replaces the old operation ‘Spell-Out’. Whereas, in the older model, Spell-Out is the point where phonetically interpretable features are stripped off the derivation and handed over to the sensory-motor interface, and narrow syntax then continues with a covert cycle to LF, there is now only a single cycle with Transfer operations to both the sensory-motor systems and the semantic systems at the same periodic points in the derivation, i.e. each cyclic boundary:



As a derivational ‘phase’, a cycle exhausts the projective potential of a head with which a given cycle sets out: if a verb, it will carry the cycle through to *v* (Aspect/Voice); if a nominal, it will carry it through to, say, D; if whatever is the head of the C-phase, perhaps T or Fin, it will be C, correlating with the computation of illocutionary Force, i.e. a full proposition put forward in discourse. Transfer thus creates a complete and then encapsulated unit of both sound and meaning. The only structure, then, that is ever computed in syntax, is of the form (37), and no syntactic object ever looks different:



Finally, he would surely not be referring simply to what (41) refers to, an unstructured and raw cat space:

(41) [cat]

In short, embedded constituents within a phase do not have full referential force: they do not determine objects of reference in the way the phase is used. They are part of an intensional description that helps to nail down a referent reached at the root. Reference is determined only at the phase level: before that, interpretation remains intensional. By the principle of compositionality, by contrast, we would be assigning independent referents to each of the constituents in (38), and within these constituents to each word. But this wouldn't explain why we only refer to one of these constituents when we use (38): why the referents of the embedded constituents are blanked out the moment we reach the root. It is a phase, therefore, which creates a unit of referential interpretation. The referent is there once the phase is complete. If it is broken off before it is maximally complete (given the projective potential of its initiating head), reference will be accordingly weaker: it may be reference to an abstract kind (cathood), a cat-space partitioned into chunks (a mass), or a cat-space partitioned into individuals: denotational specificity grows as the syntax proceeds.

Again one could, if one wanted, assign interpretations to the embedded parts of (38). The point would remain that from the viewpoint of speaker reference, their reference remains opaque. It is the phase that, by definition, determines a unit of compositional interpretation. If there were units of interpretation within a phase, there would be phases within a phase, which makes no sense: a phase is a unit of interpretation, so there are no units of interpretation within it. Compositionality in this sense fails within the phase. It makes no sense assigning independent referents to parts of phases, as compositionality does.

If phases embedded, though, compositionality might apply to their values. But strictly, as we have seen, a phase never is embedded: the complement of its head is 'transferred' the moment it is complete, and becomes inaccessible to operations then. The computation thus never sees one whole cycle embedded in another, and in this sense it never operates over a recursive structure (Arsenijevic and Hinzen 2013). Indeed, again, it makes no sense that a unit of computation embeds another unit of computation. So what gets embedded is ever only an edge, never a whole cycle. The edge corresponds to a referent, but the referent as such is an element in the discourse representation with which the syntax interfaces. It is not in the syntax. The only way the derivation can continue after a D-phase, therefore, is not by integrating the referent, John, but by integrating a *thematic role* that John *plays* with regard to an event encoded by *v*. Arguments, empirically, always play thematic roles, and are never just referential expressions. The syntax never combines referents: it creates a description to this effect. This explains why the meaning of *kill Bill* is not two referential acts juxtaposed (which it might semantically have been), one to an event of killing, one to Bill, but an event of killing of which Bill is the Theme.

Phases, then, strictly don't compose either: If the semantic value of a phase is an object, that referent does not compose or embed. What embeds is not John as such but the role of *the Agent of* a given event, *e*, say, and that description is not the same as *the Patient of* another event, *e'*, say. If embedded meanings

are not referents, two names in different embedded positions are typically not substitutable and extensionality fails. In a similar way, embedded propositions, too, will not denote truth values. They are descriptions of possible thought contents. They play a thematic role with regard to a mental state encoded in the verb taking the proposition as an argument. This is why compositionality fails. An embedded syntactic argument, *because* it is an argument, never functions as it would if it occurred alone or embedded elsewhere. It doesn't take its context-independent semantic value into the whole that contains it. It cannot take its truth value into this whole, since a truth value is a referent, and referents cannot compose. Anything that is embedded is ipso facto not maximally referentially complete. It remains intensional. If it was complete, the derivation would be cancelled: the assignment of the truth value is the stop to compositionality and recursion. As truth is the stop to compositionality, making it explicit as in any instance of (25), makes no difference to the truth-conditional content (we now understand *why* any instance of (25) holds, if it does, a fact that Tarski merely exploited). Anything that is embedded, therefore, falls short of reference to some extent; and whatever referent is determined at a phase boundary will be such that its description cannot be recovered from the derivation: by the time the referent is fixed the description is gone from the derivation.

## 7 Summary

The meaning of a sentence is not composed of independently referential parts and the way these referents are combined. By creating a part, the syntax creates an intensional description of the referent of the phasal head of which this part is a complement. It creates a whole that is *Gestalt*-like and of which the parts cannot be detached. The meaning of a part is what it is *because* the syntax has made it a part: an encapsulated unit transferred for interpretation, with only its head retained, which now becomes a description whose reference is opaque. The problem of the unity of the proposition arises only if we assume a multiplicity of independently referential parts. As for truth values, they must be unique, if syntax is bounded in the way described, and if whatever is embedded is never fully extensional. The truth value of a whole therefore never distributes over its parts, any more than the referent of a complex nominal or vP. Reference is only fixed at the left periphery of a phase. Semantic interpretation appears to wait until this happens.

Once upon a time, the phrase structure component of grammar was not recursive. A recursive phrase structure component was only introduced in the 1960s (Freidin 1999), but phrase structure has been abandoned since. Nowadays, there is only the generic operation Merge. But to understand the specific ways in which recursive structures in language arise, it is the cycle, as a mental template, that we need to understand. Recursion is mediated by phase boundaries (it is 'indirect'), and it occurs in the way it does because of these. Since a derivational phase has a semantic identity as much as a syntactic identity, recursion is semantically mediated

or conditioned. Phases however don't as such embed. Derivations consist of cycles following other cycles, not cycles embedding cycles. This coheres with the fact that the objects they compute don't embed in themselves: necessarily, they become opaque, as described. It also explains why even indirect recursion is not strict: semantic values assigned to wholes are necessarily not ones we can assign to their parts; and what parts mean when they are parts is not what they would mean if they were not parts. Syntax is for real. It affects the semantic interpretation of the parts that it connects. The idea that syntax is merely set-theoretic Merge, and that semantics consist in the composition of the referents associated to all the sets and lexical items that are merged irrespectively of where they occur, does not model this fact.

## References

- Aboh, E. (2004). Topic and focus within D. *Linguistics in The Netherlands*, 21, 1–12.
- Arsenijević, B., & Hinzen, W. (2013). On the absence of X-within-X recursion in human grammar. *Linguistic Inquiry*, 43(3), 423–440.
- Carstairs-McCarthy, A. (1999). *The origins of complex language*. Oxford: Oxford University Press.
- Chomsky, N. (1973). Conditions on transformations. In S. A. Anderson & P. Kiparsky (Eds.), *A Festschrift for Morris Halle* (pp. 232–286). New York: Holt.
- Chomsky, N. (2000). *New horizons in the study of language and mind*. Cambridge: Cambridge University Press.
- Chomsky, N. (2008a). Approaching UG from below. In U. Sauerland & H.-M. Gärtner (Eds.), *Interfaces + recursion = language? Chomsky's minimalism and the view from syntax-semantics* (pp. 1–29). Berlin/New York: Mouton de Gruyter.
- Chomsky, N. (2008b). On phases. In C. Otero (Ed.), *Foundational issues in linguistic theory*. Cambridge, MA: MIT Press.
- Chomsky, N. (2010). Some simple evo devo theses: How true might they be before language? In R. K. Larson, V. Déprez, & H. Yamakido (Eds.), *The evolution of human language: Biolinguistic perspectives* (pp. 45–62). Cambridge: Cambridge University Press.
- Cinque, G. (1999). *Adverbs and functional heads: A cross-linguistic perspective*. New York: Oxford University Press.
- Davidson, D. (2005). *Truth and predication*. Cambridge: Harvard University Press.
- De Villiers, J. (2007). The interface of language and theory of mind. *Lingua*, 117(11), 1858–1878.
- den Dikken, M. (2009). On the functional structure of locative and directional PPs. In G. Cinque & W. Schweikert (Eds.), *The cartography of syntactic structure* (Vol. 6). Oxford: Oxford University Press.
- Di Sciullo, A. M. (2005). *Asymmetry in morphology*. Cambridge, MA: The MIT Press.
- Field, H. (1994). Deflationist views of meaning and content. *Mind*, 103(411), 249–84.
- Fodor, J. A., & Lepore, E. (2002). *The compositionality papers*. Oxford: Oxford University Press.
- Freidin, R. (1999). Cyclicity and minimalism. In S. Epstein & N. Hornstein (Eds.), *Working minimalism* (pp. 95–126). Cambridge: The MIT Press.
- Hauser, M. D. (2000). *Wild minds*. New York: Holt.
- Hauser, M. D., Chomsky, N., & Fitch, T. W. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298, 1569–79.
- Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. Oxford: Blackwell.
- Higginbotham, J. (1985). On semantics. *Linguistic Inquiry*, 16, 547–931.
- Higginbotham, J. (2009). Two interfaces. In M. Piattelli-Palmarini (Ed.), *Of minds and language* (pp. 142–154). Oxford: Oxford University Press.



- Hinzen, W. (2006). *Mind design and minimal syntax*. Oxford: Oxford University Press.
- Hinzen, W. (2007). *An essay on names and truth*. Oxford: Oxford University Press.
- Hinzen, W. (2009). Hierarchy, merge, and truth. In M. Piattelli-Palmarini et al. (Eds.), *Of minds and language* (pp. 123–141). Oxford: Oxford University Press.
- Hodges, W. (2001). Formal features of compositionality. *Journal for Logic Language and Computation*, 10(1), 7–28.
- Hollebrandse, B., & Roeper, T. (2008). *Recursion as a restrictive semantic interface*. Handout from a presentation at the conference *Biosemantics*. Leiden.
- Hornstein, N. (2009). *A theory of syntax*. Cambridge: Cambridge University Press.
- Jackendoff, R. (2002). *Foundations of language*. Oxford: Oxford University Press.
- Jacobson, P. (2009). Direct compositionality. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality*. Oxford: Oxford University Press.
- Jayaseelan, K. A. (2008). Topic, focus, and adverb positions in clause structure. *Nanzan Linguistics*, 4, 43–68.
- Katz, J., & Pesetsky, D. (2009). *The recursive syntax and prosody of tonal music*. Ms., MIT.
- Larson, R. K. (2011). Clauses, propositions and phases. In A.-M. DiSciullo & C. Boeckx (Eds.), *The biolinguistic enterprise: New perspectives on the evolution and nature of the human language faculty* (pp. 366–391). Oxford: Oxford University Press.
- Lasnik, H. (2000). *Syntactic structures revisited*. Cambridge: The MIT Press.
- Leiss, E. (2009). *Sprachphilosophie*. Berlin: De Gruyter.
- Lycan, W. G. (2008). *Philosophy of language: A contemporary introduction*. New York: Routledge.
- Moro, A. (1997). *The raising of predicates: Predicative noun phrases and the theory of clause structure*. Cambridge: Cambridge University Press.
- Pawley, A. (2006). *On the argument structure of phrasal predicates in Kalam, a language of the Trans New Guinea family*. The 32nd annual meeting of the Berkeley Linguistics Society, Berkeley, CA.
- Pietroski, P. (2003). Quantification and second-order monadicity. *Philosophical Perspectives*, 17, 259–298.
- Potts, C., & Roeper, T. (2007). The narrowing acquisition path. In L. Progovac et al. (Eds.), *The syntax of nonsententials*. Amsterdam: John Benjamins.
- Prusinkiewicz, P., & Lindenmeyer, A. (1996). *The algorithmic beauty of plants* (Reprint). New York: Springer.
- Rizzi, L. (1997). The fine-structure of the left periphery. In L. Haegemann (Ed.), *Elements of grammar* (pp. 281–337). Dordrecht: Kluwer.
- Searls, D. B. (2002, November 14). The language of genes. *Nature*, 420, 211–217.
- Semenza, C. (2009). The neuropsychology of proper names. *Mind & Language*, 24(4), 347–369.
- Szabo, Z. (2009). The case for compositionality. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality*. Oxford: Oxford University Press.
- Tarski, A. (1933). Der Wahrheitsbegriff in formalisierten Sprachen. *Studia Philosophica*, 1, 261–405.
- Uriagereka, J. (2008). *Syntactic anchors*. Cambridge: Cambridge University Press.