Structural Entailment and Semantic Natural Kinds

Brendan Balcerak Jackson

(Draft: 01.05.2015)

Abstract: Is there a principled difference between entailments in natural language that are valid solely in virtue of their form or structure and those that are not? This paper advances an affirmative answer to this question, one that takes as its starting point Gareth Evans's suggestion that semantic theory aims to carve reality at the joints by uncovering the semantic natural kinds of the language. I sketch an Evans-inspired account of semantic kinds and show how it supports a principled account of structural entailment. I illustrate the account by application to a case study involving the entailment properties of adverbs; this involves developing a novel proposal about the semantics for adverbs like 'quickly' and 'slowly'. In the course of the discussion we touch on some implications of the account for the place of model-theoretic tools in natural language semantics, and about the relationship between semantic structure and logical consequence as customarily conceived.

0 Historical prologue

Richard Montague begins his landmark "English as a Formal Language" (EFL) with the crisp declaration, "I reject the contention that an important theoretical difference exists between formal and natural languages." (Montague 1974, p. 188) He goes on to demonstrate how the tools of model theory can be used to capture the way the truth conditions of sentences systematically depend on the ways they are constructed from basic lexical items for a significant fragment of English.

At the time of Montague's writings, Donald Davidson was vigorously arguing that the role of structure in determining truth conditions should be captured, not in model-theoretic terms, but by means of a Tarski-style recursive definition of truth (Davidson 1966, 1967). One important advantage of Montague's approach over Davidson's is that it supports a straightforward semantic characterization of a consequence relation for natural language: where Γ is a set of sentences and s is a sentence (all of English), s is a consequence of Γ just in case s is true in every model in which all of the sentences in Γ are true; consequence is simply the preservation of truth across all admissible variations in the interpretation of the basic lexical items.[1] This is one of the most exciting elements of EFL, because it provides an elegant way to capture relations of entailment that hold purely in virtue of the semantically relevant structures of the sentences involved. No comparable conception of structural entailment emerges naturally from the Davidsonian approach.[2]

It is thus disappointing to observe that the semantics Montague actually develops in EFL turns out not to yield *any* non-degenerate instances of structural entailment in the pure sense just defined; only reiteration – the entailment from s to s itself – comes out as structural. Non-trivial cases emerge only once Montague begins to add stipulations concerning the meanings of individual lexical items. For example, constraints on the interpretations of 'not', 'necessarily' and the 'is' of identity are added to secure some of the familiar logical consequences of sentences containing these expressions. But these are extrinsic constraints on the range of admissible interpretations that play no role at all in the account of how structure contributes to meaning,

[1] This is a simplified version of Montague's characterization. On his characterization, sentences are only true or false in a model relative to an *analysis*, which is roughly a disambiguating structural description. And Montague treats truth

[2] Davidson made efforts to construct such a conception (Davidson 1967, 1970, 1973). In my view, these efforts were ultimately unsuccessful (Balcerak Jackson 2007).

and they dilute the sense in which consequence as defined can legitimately claim to be a kind of validity in virtue of structure. It is further diluted when Montague introduces meaning postulates to capture patterns like the one exemplified by the entailment one from 'Kermit is a big frog' to 'Kermit is a frog'. This is something we can do with any entailment we like, including the one from 'Kermit is a bachelor' to 'Kermit is male'; taking this route thus amounts to giving up on trying to capture a genuine difference between validity that holds just in virtue of compositional structure and validity that also depends upon the idiosyncratic meanings of particular lexical items.

This surrender is premature. We *can* capture a genuine difference, and we can do it with essentially the same model-theoretic tools that Montague gave us. But as we will see, in order to do so we must learn to think in a very different way than Montague about how to put those tools to use.

1 Introduction

Let us consider a candidate case of structural entailment that Montague also considers in EFL:

(1) a. Rolf quickly left the stage.
b. Rolf left the stage.

The truth of (1a) guarantees the truth of (1b), and this remains the case if we replace 'quickly' with any of a wide range of adverbs: 'slowly', 'quietly', 'loudly', 'carefully', 'haphazardly', 'happily', 'angrily', and many others. It is *prima facie* plausible that the pattern in (1) ought to be captured by an adequate compositional semantics for adverbial constructions such as (1a). But then what should we say about the following?

(2) a. Rolf allegedly left the stage.
b. Rolf left the stage.

(2b) is not a consequence of (2a), since allegations that Rolf left the stage might be false. A handful of other expressions, such as 'reportedly' and 'supposedly', behave like 'allegedly' in this respect. If we hold compositional semantics responsible for explaining the entailment in (1), then the data in (2) shows that superficially similar constructions involving 'allegedly' require a different semantics. But why shouldn't we instead take the data in (2) as showing that (1) is not really structural after all? Why not give a uniform treatment of the two types of adverbs, and set aside the entailment in (1) as something that merely reflects the specific meaning of the adverb? This is the conclusion Montague reaches in EFL; it is an early instance of the now-familiar Montagovian strategy of "generalizing to the worst case." From this perspective, the entailment in (1) does not differ in kind from the 'bachelor'/'male' entailment, which no one expects to be captured by compositional semantics alone.[3]

Analogous questions arise in very many areas of semantics. For example, transitive verbs exhibit various different entailment patterns:

(3) a. Gonzo blew up the theatre.

[3] Why not trust syntactic theory to settle this question for us, by telling us whether (1a) and (2a) have the same syntactic structure or not? The problem is that even if (1a) and (2a) differ syntactically, there might be good grounds for abstracting away from this difference in the semantics. Conversely, even if (1a) and (2a) share a syntactic structure, 'quickly' and 'allegedly' might nevertheless contribute to the meaning of this structure in different ways that an adequate semantics should capture.

b. The theatre blew up.
c. Gonzo blew up.

(4) a. Gonzo ate the chicken.
b. The chicken ate.
c. Gonzo ate.

Sentence (3a) entails (3b) but not (3c); analogous behavior is observed with many other so-called *causative verbs*, such as 'boil', 'burn down', 'close', 'freeze', 'grow', and 'open'. This pattern is reversed in (4): (4a) entails (4c) but not (4b); analogous behavior is observed with many other simple activity verbs, such as 'carve', 'chop', 'drink', 'hug', 'hunt', and 'wash'.[4] Should we hold compositional semantics accountable for these patterns and develop different treatments for causatives and activity verbs (as in Parsons (1990) and Pietroski (2003))? Or should we pursue the Montagovian aim of a uniform analysis, and dismiss these patterns as due to differences in individual lexical meanings that are compositionally irrelevant?[5]

Quantificational determiners provide another example.

(5) a. All dogs bark.
b. All brown dogs bark.

(6) a. Some dogs bark.
b. Some brown dogs bark.

(5a) entails (5b), but not vice versa, and the pattern is again reversed in (6). Some determiners, such as 'no' and determiners of the form 'at most *n*', follow the pattern in (5). Others, such as those of the form 'at least *n*' and 'more than *n*', follow the pattern in (6). (Still others, such as 'most', support neither entailment.) In this case, too, capturing the entailment patterns as structural calls for different compositional treatments of constructions that are superficially similar (see e.g. Ludlow (2002)). And here, too, we can instead opt for a uniform treatment that classifies the entailments at issue as a lexical matter (as in Barwise and Cooper (1981)).

The aim of the present discussion is to develop a systematic way to address these kinds of questions. Our starting point will be a proposal that was first made, but never properly developed, by Gareth Evans. At the core of Evans's proposal is the idea that when constructing a semantic theory for a natural language, part of the goal is to discover the objectively real *semantic natural kinds* of the language. According to Evans, patterns of structural entailment help lay bare these semantic natural kinds:

> [The structural/non-structural contrast is] a contrast between inferences whose validity depends merely upon the *kind* of semantic elements out of which a sentence is constructed, and its manner of construction, on the one hand, and inferences whose validity depends upon the specific variation a particular semantic element is playing upon the theme all expressions of its kind must play, on the other (Evans, 1985, pp. 60-61).

From this perspective, we can see the task of accounting for patterns of structural entailment as a central part of the effort to construct a theory that succeeds in carving semantic reality at its joints. Conversely, the question of how to classify a given pattern can be approached by

[4] For a rich survey of these and many other types of verb alternations see Levin (1993).
[5] In some languages there are overt morphological differences between the causative and the inchoative uses of verbs such as 'blow up' (illustrated in (3a) and (3b), respectively). But this does not automatically settle the issue of how the English constructions, which do not exhibit such morphological differences, ought to be handled.

investigating how the semantic natural kinds in the relevant area are to be individuated. Do 'quickly' and 'allegedly' belong to a single semantic natural kind or to two distinct kinds? Do the patterns in (5) and (6) correspond to distinct species of quantificational determiners? The answers to these kinds of questions are not immediately obvious, to be sure, but we have no reason to doubt that they *have* substantive answers – at least, none that does not derive from a more sweeping conventionalism about natural kinds in general.

To make good on Evans's proposal, we need to develop some positive conception of semantic natural kinds. What determines the natural kind to which a given lexical item belongs? And how exactly do semantic kinds, so conceived, figure into a general characterization of structural entailment? Answers to these questions are developed in §2 and §3. What emerges from the discussion in these sections is not a radically new way of doing semantics, but a certain foundational conception of what is involved in doing semantics as it is (at least often) already done. As we will see, there is value in making this conception explicit and employing it deliberately. To that end, §4 applies the conception developed in §2 and §3 to the case with which we began, by looking in some detail at adverbs like 'quickly' and 'allegedly'. I argue that there is very strong evidence for classifying 'quickly'-type adverbs as belonging to a distinct natural kind that excludes 'allegedly'-type adverbs, and I sketch a novel semantics for the former that vindicates the claim that the pattern in (1) is genuinely structural. In §5 we conclude by briefly returning to Montague and EFL.

Two foundational assumptions will be important in what follows. The first assumption is *compositional realism*, according to which the semantic properties of simple lexical items are explanatorily prior to the semantic properties of complex expressions, and of sentences in particular. It is in part *because* a given sentence *s* is constructed out of lexical items with certain semantic properties that *s* has the particular meaning it has; the semantic properties of those lexical items are prior to and help ground the fact that *s* has that meaning. This assumption runs counter to deflationary accounts that see lexical items as bearing only trivial semantic relations to the complex expressions in which they occur (Horwich 2001). It also runs counter to instrumentalist views that see no basis in semantic reality for choosing between a range of non-equivalent compositional theories that yield the same results at the level of sentence meanings (see e.g. Davidson (1967) and (1970), Lewis (1975), and Wright (1981)). Compositional realism maintains that there is an inquiry-transcendent domain of substantive facts about the semantic properties of lexical items and how they contribute to the semantic properties of complex expressions.[6]

The second assumption is a claim about what sorts of semantic properties these are. In what follows we will assume that semantics aims to explain how the *truth conditions* of sentences of the language under study depend on their parts and the way they are put together. (Many, if not all natural language sentences have truth conditions only relative to a given context of use; this complication will be suppressed for most of the present discussion.) Other conceptions are available, of course. For example, one might hold that the job of compositional semantics is to map sentences onto propositions, conceived either as structured entities composed of individuals, properties and relations (King 2007, Salmon 1986, Soames 1987) or as something like Fregean thoughts (Bealer 1982, Chalmers 2002). I will not try to justify the truth-conditional conception here, other than to observe that the alternatives yield no explanation for entailment until they are supplemented with an account of truth conditions for propositions; such an

[6] Szabó (in preparation) defends a qualified version of compositional realism, according to which the semantic properties of *open category* lexical items (e.g. paradigmatic nouns, verbs, adjectives and adverbs) but not of *closed category* lexical items (e.g. conjunctions, modal auxiliaries and determiners) are explanatorily prior to the semantic properties of the complex expressions containing them. Here I make no use of the distinction between open and closed categories.

account is likely to recapitulate much of the work that goes into developing a semantics directly in terms of truth conditions.[7]

There is little question that both of these foundational assumptions lie in the background of Evans's proposal, and they will also inform the account of semantic natural kinds and structural entailment developed here.

2 Semantic natural kinds as alethic essences

In virtue of what does a given lexical item *e* belong to a given semantic kind *K*? Evans's original discussion serves as our starting point.

According to Evans, we should think of semantic natural kinds as determining:

> … for each kind of semantic expression, an entity – a set, a truth value, a function from sets to truth values, or whatever – which may appropriately be assigned to members of that kind upon an arbitrary interpretation of the language. We can regard the specification of the kind of assignment as a specification of the underlying real essence which a word has in common with many other words, and of which the validity of certain inferences involving it is a consequence. These will be the structural valid inferences, inferences which are truth-preserving no matter how we permute assignments within the limits laid down as appropriate for members of that category. (Evans 1985, p. 61)

Kent Johnson (2003) reads Evans as endorsing what we might call a *nominalist* approach. According to Johnson, Evans is committed to the following:

> … (i) two semantic categories are identical if and only if they are associated with the same logical entity; (ii) two logical entities are identical if and only if they specify the same set of inferential properties, and (iii) a word belongs to a given semantic category only if the word has all the logical properties that are associated with the category. (Johnson 2003, p. 178)

As Johnson emphasizes, the "logical entities" in this account are ultimately superfluous: clauses (i) and (ii) associate each semantic kind with a set of inferential properties, and clause (iii) specifies that the semantic kind membership of a lexical item ultimately depends on the inferential properties it possesses. This contrasts with a *realist* approach, on which the fact that *e* has certain inferential properties is taken to be a symptom of its possession of a certain underlying "semantic essence" that helps give rise to the valid (and invalid) entailments in which *e* is involved. Inferential properties thus play an evidential role rather than a constitutive role for the realist; it is the underlying semantic essences that determine semantic kind membership.

Johnson offers two considerations in support of the nominalist reading of Evans. The first is that it gives us a very straightforward way of using semantic kinds to answer our original question about how to classify entailments as structural or non-structural: an entailment is structural just in case the corresponding inferential property is a member of a set that turns out to individuate a semantic kind to which one or more of the lexical items occurring in it belongs. Johnson's second consideration is that, "if we did not identify a word's semantic essence with its inferential

[7] An analogous point applies to internalist approaches that see compositional semantics as mapping expressions onto structured mental representations (e.g. Jackendoff 1990, McGilvray 1998). Some internalists would deny that a semantic theory ought to offer explanations of entailment in the first place; perhaps semantics should instead aim to account for certain transitions in the reasoning of competent speakers (Kempson 2011). I will not try to enter into this debate here.

properties, we would have to identify this essence with some other thing E. But it is entirely unclear what E could be, and it is unclear what it would be for the inferential properties to 'flow' from E." (Johnson 2003, p. 178)

The nominalist reading of Evans does not survive scrutiny, however. One problem is that Evans quite clearly intends semantic kinds to help *explain* the validity of certain entailment patterns. But it is not at all obvious that nominalist kinds can play this role: it is not very illuminating to say that a lexical item licenses a certain entailment because it belongs to kind *K*, when belonging to *K* simply consists, in part, in being such as to license that entailment. This is like being told that some sample substance conducts electricity because it is a metal, only then to be told that for something to be a metal *just is* in part for it to conduct electricity.

A related problem is that if we determine kind membership in terms of inferential properties, then we simply trade the question of whether a given entailment counts as genuinely structural for the question of whether the corresponding inferential property is one that matters for kind membership. Determining whether 'quickly' and 'allegedly' belong to the same semantic kind, for example, immediately leads to the question of whether or not the entailment in (1) is one that matters for determining semantic kind membership. Our hope was to get a better grip on the status of the entailment by recourse to semantic natural kinds. But on the nominalist approach, this recourse simply brings us back to the status of the entailment again.

In fact, this is a problem that Evans himself notes immediately after the passage quoted above. He goes on to address as follows:

> What makes a parallel problem in the taxonomy of natural kinds tractable is the conception of kinds as being differentiated by underlying structures from which the characteristics used in classification may be regarded as flowing. To solve our problem for semantic kinds we need to find room for a parallel conception of something from which an expression's inferential properties may be regarded as flowing. (Evans 1985, p. 61)

Here Evans clearly identifies the realist's semantic essences ("underlying structures") as the solution to the difficulty the nominalist faces when trying to individuate semantic kinds in terms of inferential properties. For Evans, inferential properties clearly play a merely evidential role in helping to identify these essences; they are the characteristics "used in classification." It is very difficult to read this passage as recommending anything other than a realist approach.

What then about Johnson's worry that the realist's appeal to underlying semantic essences is mysterious? In fact, the mystery is greatly overstated. There is no need to invoke some unfamiliar class of *sui generis* properties. The sorts of properties already invoked in ordinary empirical work in truth-conditional semantics – properties such as having a certain reference, standing for a certain property or relation, being satisfied by entities that meet such-and-such conditions, and more generally properties that are taken to figure into the determination of truth conditions in one way or another – are perfectly well suited to play the role of the underlying essences that determine kind membership. And there is nothing mysterious about the idea of entailments "flowing from" such properties, especially given the foundational assumptions articulated in §1. A given sentence *s* has the particular truth conditions it does partly because it contains lexical items with certain truth-conditional properties, and when *s* entails some other sentence *s´* this, too, is partly explained by the truth-conditional properties of their constituent lexical items. This is true whether the entailment at issue is ultimately to be classified as structural or not. For someone who thinks of semantics in this way, as Evans certainly does, Johnson's mysteriousness worry is bound to have very little force.

On the whole, then, Evans is most plausibly read as recommending a realist rather than a nominalist approach to semantic kinds. If we follow Evans's recommendation, where might this lead?

We can begin by noting that truth-conditional properties come in various degrees of specificity and stand in determinate-determinable relations. For example, the property of referring to Kermit is a determinate instance of the more general property of referring to an individual, and the property of being satisfied by male persons is a determinate instance of the more general property of being satisfied by persons, which is itself a determinate instance of the property of being satisfied by individuals. Suppose that lexical items e_1, …, e_n each make their own distinctive contribution to truth conditions by virtue of possessing determinate truth-conditional properties P_1, …, P_n, respectively. They might nevertheless all contribute *in the same way* to truth conditions by virtue of the fact that P_1, …, P_n are all instances of some determinable property P. It is natural to appeal to determinate-determinable relations of this sort to characterize semantic kind membership: each semantic kind *K* is associated with some determinable truth-conditional property P, and for *e* to belong to *K* is simply for *e* to instantiate some determinate instance of P. On this view, semantic kinds correspond to distinctive ways of contributing to truth conditions, and the specific contribution *e* makes is its determinate instance of the more general way expressions of that kind contribute – it is *e*'s "particular variation on the theme all expressions of its kind must play," as Evans puts it.[8]

We need not and should not allow that everything we can describe as a way of contributing to truth conditions corresponds to a genuine semantic kind; semantic kinds are discovered, not invented or stipulated. For example, expressions like 'bachelor', 'husband' and 'brother' all have in common that they apply only to males. But we can determine whether a property such as *being satisfied by male individuals* carves out a kind only by determining what kind of explanatory work, if any, the property is called on to do within an adequate compositional semantics. Let us call the determinable truth-conditional properties that determine genuine semantic kinds *alethic essences*. Classifying a given property as an alethic essence – and hence classifying the lexical items that possess it as belonging to a genuine natural kind – calls for detailed investigation of specific linguistic phenomena, and evaluation of the explanatory costs and benefits of competing hypotheses to account for those phenomena.

To help further elucidate the realist account, in the remainder of this section we discuss three further features that distinguish it from – and make it more attractive than – the nominalist alternative with which we began.

First feature: non-circularity. The first advantage of the realist account has, in effect, already been noted: it yields an informative characterization of structural entailment patterns as cases in which validity is explained by the way the sentences involved are constructed out of lexical items of certain semantic kinds. On the realist approach, this amounts to seeing structural entailments as cases in which validity can be fully explained in terms of the general kind-individuating features of lexical items, and the way such kind-individuating features contribute to the truth conditions of the structures in which they occur. On the nominalist approach, by contrast, semantic kinds at best only yield explanations of the *virtus dormitiva* sort.

[8] Szabó (in preparation) identifies certain determinable truth-conditional properties that he takes to determine whether a given lexical item counts as a noun, verb, adjective or adverb. From the realist perspective articulated here, this amounts to a treatment of these categories as highly general semantic natural kinds. Note that Szabó's proposal leaves open the possibility that the lexical items in each category can be sorted into various further theoretically important sub-kinds, sub-sub-kinds and so on.

Second feature: diversity at the nominal level. For the realist, inferential properties are treated as symptoms of the underlying truth-conditional properties in terms of which semantic kinds are individuated. There is no reason to think that *only* inferential properties can play this role. For one thing, the distributional behavior of a lexical item can help shed light on the semantic kind to which it belongs, since the truth-conditional properties of a lexical item help determine which constructions involving it are acceptable and which are marked. For example:

(7) a. Scooter is aware that Fozzy is on stage.
b. * Scooter is likely that Fozzy is on stage.

(8) a. It is likely that Fozzy is on stage.
b. * It is aware that Fozzy is on stage.

The data in (7) and (8) are plausibly best explained by the hypothesis that the former allows a subject argument (and in fact requires one) while the latter does not (Williams 1983); this amounts to an explanation of the distributional data in (7) and (8) in terms of the underlying truth-conditional properties of 'aware' and 'likely'. To take another example, there are many well-known distributional differences between adjectives like 'green' and adjectives like 'fake', such as the ability of the former, but not the latter, to occur with degree modifiers and in comparative constructions:

(9) a. a very green frog
b. * a very fake frog

(10) a. Kermit is more green than Robin.
b. * Kermit is more fake than Robin.

There is wide consensus that an adequate semantics for modification by words like 'green' and 'fake' should help explain these differences. Even patterns of distributional behavior that are ultimately best explained in purely syntactic terms can be a useful guide when identifying semantic kinds, since the way a lexical item contributes to truth conditions surely has something to do with the positions it can occupy in various syntactic configurations.

Patterns of ambiguity and non-ambiguity are further symptoms of underlying truth-conditional similarities and differences. For example, consider the following:

(11) a. Janice quickly taught Floyd how to play the guitar.
b. Floyd was quickly taught by Janice how to play the guitar.

(12) a. Janice reluctantly taught Floyd how to play the guitar.
b. Floyd was reluctantly taught by Janice how to play the guitar.

The sentences in (11) are unambiguous and equivalent. But (12b) is ambiguous: it can be read as attributing reluctance either to Janice or to Floyd. This reveals a difference in underlying truth-conditional properties – 'reluctantly' is semantically sensitive to the arguments of the verb it modifies in a way that 'quickly' is not – and such differences might turn out to be important for sorting adverbs like 'quickly' and 'reluctantly' into natural kinds (McConnell-Ginet 1982).

The general methodological point is that the realist need not to try to circumscribe in advance the kinds of phenomena at the "nominal level" that might turn out to play an important evidential role in taxonomizing lexical items. Inferential, distributional and ambiguity phenomena certainly play a role, as we have just seen. But other phenomena such as presuppositions, generalized

implicatures and other kinds of discourse behavior might also shed light on underlying alethic essences in some cases. Further foundational commitments might also recommend additional sources of evidence. For example, if we assume that the properties and operations we posit in our semantics must somehow be cognitively accessible to and utilized by competent speakers in language comprehension, then we might be led to take into account psychological and neurological evidence of various kinds (see e.g. the work of Pietroski et al. (2009) and Lidz et al. (2011) on 'most'). However the nominal level is circumscribed, the more numerous and the more varied the commonalities are that we find at this level, the more reason we have to search for some shared underlying truth-conditional properties that contribute to a unifying explanation of them. To discover these properties is to discover the underlying real essence that individuates the kind.[9]

Third feature: cross-linguistic applicability. It is extremely plausible that many, if not all, semantic natural kinds have instances in more than one language; at the very least, this is surely an empirical hypothesis that is open to cross-linguistic investigation. On the nominalist approach, however, it is not at all clear whether there is any plausible way to make sense of cross-linguistic claims about semantic kinds. For example, we might want to know whether German contains lexical items of the same kind *K* as 'quickly', and if so which these are. For the nominalist, membership in *K* is determined by looking at inferential properties; let us suppose that the property corresponding to (1) is one of them. Does the German lexical item 'schnell' share this property? We cannot simply check by replacing 'quickly' with 'schnell' in (1); the result is meaningless word salad rather than a valid (or invalid) entailment. What the nominalist needs is some way to identify certain German constructions as the relevant analogs of the English entailments that are used to determine membership in *K*. The problem is that identifying such constructions inevitably leads us to consider German sentences such as the following:

(13) Rolf hat die Bühne angeblich verlassen.
Rolf-*nom* the stage-*acc* allegedly leave-*past*
Rolf allegedly left the stage.

Is (13) analogous to 'Rolf quickly left the stage' in the relevant sense? It is not at all clear how the nominalist would have us decide this. Moreover, if we *did* have a way of deciding then it could presumably also be applied directly to the question of whether 'Rolf quickly left the stage' is of the same relevant kind as 'Rolf allegedly left the stage'. This would give us an independent way to determine whether the pattern in (1) counts as structural, and semantic kinds as the nominalist conceives them would be rendered superfluous. This underscores again just how little the nominalist's kinds really contribute to an informative account of structural entailment.

In fact, an analogous problem for the nominalist arises even within a single language, because of the possibility of sub-categorization restriction violations and the like.

(14) a. Rolf quickly discovered the proof; so Rolf discovered the proof.
b. ?? Rolf loudly discovered the proof; so Rolf discovered the proof.

(15) a. Kermit quietly slept; so Kermit slept.
b. ?? Kermit urgently slept; so Kermit slept.

[9] It is useful to compare the present view of semantic kinds to one that treats them as *homeostatic property cluster* kinds, i.e. as clusters of properties whose regular co-instantiation is explained by the operation of homeostatic mechanisms (Boyd 1999, 1989). The homeostatic property cluster view is best understood as a version of the nominalist approach, one that identifies semantic kinds with clusters of distributional, entailment and other properties at the nominal level; alethic essences figure into the linguistic processes that play the role of the mechanisms responsible for the clustering. On the present view, by contrast, semantic kind membership is simply determined directly by the alethic essences.

Suppose that the entailments in (14a) and (15a) turn out to be valid in virtue of the semantic kind to which 'quickly' and 'quietly' belong. The infelicity of (14b) and (15b) should not *automatically* rule out 'loudly' and 'urgently' as belonging to this kind as well. As in the cross-linguistic case, the nominalist needs some way to make sense of the hypothesis that 'loudly' and 'urgently' participate in the same *kinds* of entailments as 'quickly' and 'quietly' even if they cannot be everywhere uniformly substituted for one another. And again, the danger is that once this need is met the nominalist's semantic kinds become superfluous.

On the realist account cross-linguistic claims about semantic kinds are unproblematic: alethic essences are not language-specific, and any two lexical items that have the same alethic essence are members of the same kind, regardless of the languages to which they belong. Likewise, 'quickly', 'loudly' and 'urgently' can share an alethic essence even if they possess different determinate instances of it, and these determinate differences result in different restrictions on the verbs they can be used to modify. More generally, a lexical item might belong to a given kind even when it fails to exhibit one or more of the diagnostic criteria that are associated with the kind, for there might be plausible independent explanations for why the alethic essence the lexical item shares with other members of the kind fails to manifest its usual symptoms in certain circumstances.

A further important and attractive feature of the realist approach, and one that will loom large in the discussion to follow, is that it provides us with a fruitful way to think about the application of model-theoretic tools to the semantics of natural language. The next section examines this feature in more detail.

3 Interpretational semantics and categorial validity

Let an *interpretation* **I** for a language *L* be a function that assigns to each basic lexical item of *L* some entity, paradigmatically an extra-linguistic entity such as an individual, a set, or a function of some sort.[10] An *interpretational semantics* for *L* has two components: a specification of what counts as an *admissible* interpretation of the basic expressions of *L*, and a set of *composition rules* that together specify interpretation-relative truth conditions for each sentence *s* of *L*. (We continue to ignore the role of extra-linguistic context.)

Let us call an interpretational semantics *Montagovian* when it constructs these two components on the basis of a theory of types. A theory of types begins with a specification of certain classes of entities as basic types; in ELF, for example, the basic types are **e**, the class of actual and possible individuals, and **t**, the class of truth values {0,1}. Further types are then defined in terms of recursive operations on types. In ELF, for example, whenever *x* and *y* are types then the class of functions with domain in *x* and range in *y*, denoted <*x*,*y*>, is also a type. (Here we follow contemporary notational conventions rather than the ones used in EFL.) Associating each lexical item with a type allows for a neat characterization of admissible interpretations: an interpretation **I** is admissible just in case the value **I** assigns to each lexical item *e* is drawn from *e*'s associated type. Types also guide the formulation of composition rules: the rules governing a given construction must be such that they can apply to the lexical items it contains no matter how they are interpreted within the constraints imposed by their associated types.

From the perspective of the realist theory of kinds developed in §2, it is very natural to see a Montagovian interpretational semantics as providing us with a way to formally represent

[10] I also sometimes refer to **I**(*e*) as the interpretation of *e* relative to **I**, and rely on context to resolve the ambiguity.

hypotheses about semantic natural kinds. When an interpretation **I** assigns a certain value to a given lexical item *e*, this can be seen as a way of representing a determinate truth-conditional property of *e* in virtue of which **I** takes it to make its particular contribution to truth conditions; if **I**(*e*) is some individual *x*, for example, then **I** treats *e* as having the property of referring to *x*. Likewise, when we associate *e* with a certain type we represent it as having a certain determinable truth-conditional property, one whose possible determinate instances are represented by the various interpretations permitted by *e*'s associated type; if *e* is associated with **e**, for example, then an interpretation is admissible only if it treats *e* as referring to some individual or other. This determinable truth-conditional property is one that *e* shares with all other lexical items associated with the same type, and it is a property that *e* must have if it is to play the compositional role that the rules of the theory ascribe to it. In short, it is what the theory takes to be *e*'s alethic essence, the property that makes *e* the kind of lexical item it is.

A Montagovian characterization of admissible interpretations thus embodies a proposal about how to individuate the semantic kinds of the language, and the compositional component of the theory embodies a proposal about how the kinds identified contribute to truth conditions.[11] Of course, Montague himself almost surely did not think of things this way. For one thing, notice that on the present conception, any reason we have for thinking that two superficially similar lexical items in fact belong to distinct semantic kinds is a reason for associating them with distinct types that impose distinct constraints on their admissible interpretations. This, in turn, might call for adjustments in the way truth conditions are determined for constructions containing the two kinds. This means that investigating where and how to draw the boundaries between semantic kinds can give us good reason to develop different compositional treatments of superficially similar constructions, even when a uniform treatment is, in principle, available. In fact, as we will see in §4, this is exactly what we find when we turn to the semantics of adverbs like 'quickly' and 'allegedly'. Uniform treatments are not a *desiderata* for a Montagovian semantics viewed as a way of carving semantic reality at the joints.

Let us turn next to structural entailment. Where **T** is an interpretational semantics for *L* and *s* and *s´* are sentences of *L*, we can define *categorical validity* as follows:

> The entailment from *s* to *s´* is categorially valid relative to **T** iff *s´* is true relative to every interpretation admissible on **T** relative to which *s* is true.[12]

The entailments that qualify as categorially valid relative to **T** are the ones that depend only on the composition rules of **T** and whatever features of lexical items are preserved by all interpretations defined as admissible by **T**; any entailments that depend on the specific meanings of certain lexical items are washed out by the variations across interpretations. Since the limits on these variations embody the theory's hypotheses about semantic kinds, there is a clear sense in which the entailments that are categorially valid according to **T** are the ones whose validity we can explain just in terms of the ways the sentences are constructed out of lexical items of certain natural kinds. It is just these entailments that Evans proposes to classify as the ones that **T** treats as structural. Within the interpretational framework, then, a division between the structural

[11] One consequence of this way of thinking is that only simple lexical items enter into semantic kind classifications; complex expressions fall outside the domain of the interpretation functions that embody a theory's claims about semantic kinds. This is not to deny that some class of complex expressions might turn out to play the same kind of role in determining truth conditions as lexical items of a certain natural kind. (This might be the case, for example, with some adverbs and complex adjuncts; see Parsons 1990, chapter 4.) By analogy, suppose that a species of symbiotic bacteria plays a certain role in, say, the digestion process of a larger organism; perhaps the role is to produce a certain enzyme. If we later find a complex suite of interacting bacteria that plays exactly the same role, we could acknowledge this even without counting the complex suite of bacteria as a member of our original species.

[12] This definition can be extended in the obvious way to apply to sequences of premises and conclusions of any countable length.

entailment patterns and the rest emerges naturally out of the process of developing an adequate taxonomy of semantic natural kinds of the language.

Notice that no assumption has been made that lexical items such as 'and', 'not' and 'every' are syncategorematic expressions that fall outside the domain of interpretation functions, or that they receive fixed interpretations as in EFL. This implies that there is no built-in guarantee that the entailments customarily classified as logical come out as categorially valid. This is as it should be: whether logical entailment patterns turn out to be structural or not – and if so, which ones – depends on where the logical constants turn out to fit in the best overall taxonomy of semantic natural kinds. Are 'and' and 'or' distinct natural kinds of connectives? Are 'every' and 'some' distinct natural kinds of quantifiers? The answers to these sorts of questions should not be judged in advance merely by our choice of semantic frameworks.[13]

We conclude this section with a brief remark about meaning postulates, of which no mention has so far been made. Indeed, there is no natural role for them within the present conception. The effect of a meaning postulate is to impose further constraints on admissible interpretations by stipulating that every interpretation render certain sentences true, or certain entailments truth preserving. In this way meaning postulates provide an additional tool for securing categorial validity. But if a given meaning postulate is meant to capture an entailment pattern involving every lexical item of a certain natural kind, then the validity of the pattern ought instead to be captured by articulating appropriate interpretations for lexical items of that kind. Only then can we claim to have identified the alethic essences that ultimately help ground the entailment pattern, and thus to have shown why the entailment counts as genuinely structural. On the other hand, if the postulate singles out specific members within a certain kind, or cuts across kinds, then the entailments it secures are ones that should *not* be rendered categorially valid by a theory that aims to carve semantic reality at the joints; in this case the postulate merely serves to muddle distinctions that the theory has a responsibility to make clear. On the present conception, then, meaning postulates have no worthwhile contribution to make to the explanation of structural entailment.[14]

4 Natural kinds of adverbs

Let us return to the entailment involving 'quickly' with which we began in §1:

(16) a. Rolf quickly left the stage.
b. Rolf left the stage.

Call this the *factive* entailment pattern. Is the factive pattern genuinely structural?

As noted in §1, the factive pattern is supported by all so-called *manner adverbs*, such as 'slowly', 'loudly', 'smoothly', 'roughly', 'gently', 'forcefully', 'carefully', and 'sloppily'. Early accounts of manner adverbial modification are versions of the *operator approach*, which represent the meanings of manner adverbs as functions from verb-phrase meanings to verb-phrase meanings. On this

[13] One source of the temptation to assume that logical entailments are structural is the historically influential idea that logical constants do not function as meaningful lexical items, but rather merely serve somehow to reflect differences in syntactic structure. I raise some problems for this assumption elsewhere (Balcerak Jackson 2006).

[14] In addition, Zimmermann (1999) shows that in many cases it is difficult to formulate meaning postulates generally enough to capture the desired entailments while avoiding inconsistency. I agree with Zimmermann (1999, p. 559) that "meaning postulates should not be regarded as proper tools of natural language semantics." But it should be clear from the discussion here that I disagree with Zimmermann's subsequent conclusion that model-theoretic semantics is ultimately dispensable as well.

approach the meaning of 'quickly' combines with the meaning of 'leave the stage' to yield a new meaning of the same type (Cresswell 1974, Montague 1974, Thomason and Stalnaker 1973).[15]

The operator approach easily explains one of the basic semantic features of manner adverbial constructions, their *non-intersectivity*. For example, suppose that Scooter reached the rafters by climbing a rope:

(17) a. Scooter quickly climbed the rope.
b. Scooter did not quickly reach the rafters.

The truth of (17a) is consistent with the truth of (17b), even if Scooter's action of climbing is (or constitutes) his action of rafter reaching; his action might be quick for a rope climbing and yet slow for a rafter reaching. (Perhaps people usually just take the elevator.) On the operator approach this is due to the fact that 'quickly' is sensitive to the meanings of the verb phrases on which it operates, and 'climb the rope' and 'reach the rafters' are not synonymous. The operator approach also generalizes straightforwardly to modifiers such as 'allegedly' and 'supposedly'; as we have seen, this is why Montague takes an operator approach in EFL.[16] However, the trade-off is that there is no way for the approach to capture the factive entailment pattern as categorially valid, except by *ad hoc* stipulations on the functions that can be assigned to 'quickly' and other manner adverbs.

From the present perspective, the operator approach embodies a treatment of manner adverbs and adverbs like 'allegedly' as lexical items belonging to a single natural kind. Within such a theory we might still hope to capture the factive pattern via meaning postulates, as both Montague in EFL and Thomason and Stalnaker (1973) suggest. But as we noted in §3, this would only serve to obscure the fact that the operator approach is committed to a taxonomy according to which the pattern does not count as genuinely structural; it treats the pattern merely as an expression of idiosyncratic features of the specific meanings of some (but not all) instances of the common kind to which 'quickly' and 'allegedly' belong.

However, there is very good reason to conclude that this taxonomy is not adequate to the facts. Manner adverbs exhibit a rich array of shared behavior over and above factivity, and their behavior is in many ways quite different than that of adverbs like 'allegedly' and 'reportedly'. To begin with, there are further differences in inferential properties:

(18) a. Rolf quickly left.
b. Rolf's leaving was quick.
c. Quickly was how Rolf left.

(19) a. Rolf reportedly left.
b. Rolf's leaving was reported.
c. ?? Reportedly was how Rolf left.

The (a) sentence in (18) entails the (b) and (c) sentences. But (19a) does not entail (19b) or (19c), both of which require that there was an act of leaving to report; in addition, (19c) is highly marked. Manner adverbs also *fail* to support certain entailments that are supported by 'allegedly'-type adverbs, such as the following:

[15] The operator approach is compatible with different views about the meanings of verb phrases over which such operators are defined.
[16] Thomason and Stalnaker (1973) reject Montague's treatment of 'allegedly' in particular as a verb phrase operator, but they preserve the operator approach to manner adverbs on the grounds that it can be extended to modifiers like 'partially' and 'halfway' (as in 'Kermit partially filled the bucket').

(20) a. Rolf allegedly left.
b. It is allegedly true that Rolf left.

(21) a. Rolf quickly left.
b. ?? It is quickly true that Rolf left.

Another noteworthy difference in entailments is illustrated by the following:

(22) a. Gonzo gently kissed Camilla.
b. It is not the case that Gonzo gently kissed a chicken.

(23) a. Gonzo allegedly kissed Camilla.
b. It is not the case that Gonzo allegedly kissed a chicken.

The sentences in (22) together imply that Camilla is not a chicken. But this is not implied by the sentences in (23), since the allegations that Gonzo kissed Camilla might never have mentioned that she is a chicken.[17] The operator approach can accommodate these further patterns of entailment and non-entailment only by additional stipulations, and it has no natural explanation for why the different patterns cluster together in this way.

There are also several distributional differences between the two classes of adverbs. One telling difference is that 'allegedly' can occur with any sort of verb phrase, as well as with non-verbal predicates, while manner adverbs are restricted to verb phrases like 'tell a joke' that describe events rather than enduring states or properties of individuals:

(24) a. Fozzy supposedly told a funny joke.
b. Gonzo allegedly knows French.
c. Kermit is reportedly allergic to peanuts.

(25) a. Fozzy loudly told a funny joke.
b. * Gonzo quietly knows French.
c. * Kermit is quickly allergic to peanuts.

If manner adverbs function as verb-phrase operators there is no obvious reason why their application should be restricted in ways that verb-phrase operators in general are not. Notice, moreover, that the restriction is not predicated by the constraints on lexical meaning that the operator account must posit to capture factivity; the operator account must again invoke further independent constraints. But then it is puzzling that we do not find any adverbs subject to the complementary constraint, i.e. adverbs that can only occur with predicates that express states or properties and not with verbs that express events (Katz 2003).

There are also systematic restrictions on the positions manner adverbs can occupy in the sentence, relative to other elements such as negation and adverbs of quantification like 'often':

(26) a. Scooter often quietly leaves by the backstage door.
b. ?? Scooter quietly often leaves by the backstage door.

[17] This observation is made in Eckardt (1998), who attributes it to Ede Zimmermann. The data is made slightly more complicated by the fact that there is arguably a *de re* reading of (23b) according to which it is not the case that there is some chicken *x* such that Gonzo has been alleged to have kissed *x*. On this reading, the sentences in (23) do together imply that Camilla is not a chicken. The pertinent difference, then, is that there is *no* reading of the sentences in (22) on which they are jointly consistent with Camilla's being a chicken.

(27) a. Scooter didn't quietly leave by the backstage door.
b. ?? Scooter quietly didn't leave by the backstage door.
c. ?? Quietly, Scooter didn't leave by the backstage door.

Adverbs like 'allegedly' are subject to no such restrictions:

(28) a. Scooter often allegedly leaves by the backstage door.
b. Scooter allegedly often leaves by the backstage door.

(28) a. Scooter didn't reportedly leave by the backstage door.
b. Scooter reportedly didn't leave by the backstage door.
c. Reportedly, Scooter didn't leave by the backstage door.

At the same time, manner adverbs are more flexible in their occurrence relative to many other adjuncts than 'allegedly'-type adverbs:

(29) a. Piggy forcefully hit Kermit in the head with a frying pan.
b. Piggy hit Kermit forcefully in the head with a frying pan.
c. Piggy hit Kermit in the head forcefully with a frying pan.

(30) a. Piggy allegedly hit Kermit in the head with a frying pan.
b. ? Piggy hit Kermit allegedly in the head with a frying pan.
c. ? Piggy hit Kermit in the head allegedly with a frying pan.

The sentences in (29) are all intuitively equivalent. But the sentences in (30), to the extent that they are interpretable, intuitively mean different things: (30a) is neutral on whether any hitting actually occurred; (30b) suggests that a hitting definitely occurred, while it was only alleged that it was in the head with a frying pan, while (30c) suggests that a hitting occurred and that it was definitely in the head, while it was only alleged that it was with a frying pan.

What all this illustrates is that the factive entailment pattern is just one element in a rich and diverse array of behavioral patterns that manner adverbs share, and that they do not share with adverbs like 'allegedly' and 'supposedly'. The uniform treatment of both classes of expressions as verb-phrase operators cannot account for these patterns except by adding more and more *ad hoc*, unrelated stipulations and constraints; it has no genuine unifying explanations to offer. The fact that it can offer a uniform compositional treatment of 'Rolf quickly left the stage' and Rolf allegedly left the stage' is a rather poor trade-off. It is much more plausible simply to concede that the operator approach fails to cut semantic reality at the joints.

If manner adverbs constitute a natural kind that excludes 'allegedly', what is the underlying alethic essence that individuates the kind? We can start by observing that there is surely a close semantic relationship between adverbs like 'quickly' and 'slowly' and adjectives like 'quick' and 'slow' that have uses as event-level adjectives modifying event-denoting noun phrases ('game', 'explosion') and gerundives ('stabbing', 'leaving'). Like ordinary individual-level adjectives, event-level adjectives have attributive as well as predicative uses:

(31) a. The game we attended yesterday was <u>slow</u>.
b. The forensic evidence indicated that the stabbing was <u>brutal</u>.

(32) a. Waldorf and Statler gave each other a <u>quick hug</u>.
b. The forensic evidence indicated that it was a particularly <u>brutal stabbing</u>.

On a standard neo-Davidsonian event-based semantics, many verbs also behave as predicates of events, analogous 'game' and 'stabbing', so that a sentence like (33a) has truth conditions that can be perspicuously represented as in (33b).[18]

(33) a. Waldorf hugged Statler.
b. ∃*e* (Hug(*e*) & Agent(*e*,*waldorf*) & Theme(*e*,*statler*))

This suggests that from the point of view of semantics, modification by a manner adverb is analogous to the attributive use of an event-level adjective. If so then 'quickly' in (16a) has essentially the same semantic function as 'quick' as it occurs in the following:

(34) There was a quick leaving of the stage by Rolf.

Many neo-Davidsonians have developed accounts along these lines (Eckardt 1998, Higginbotham 1985, Katz 2003, McConnell-Ginet 1982, and Parsons, 1990).

On this approach the inability of manner adverbs to occur with predicates such as 'knows French' and 'allergic to peanuts' is to be expected, since the same phenomenon occurs in the adjectival case:

(35) a. ?? Gonzo's quiet knowledge of French
b. * Kermit's quick allergy to peanuts

The noun phrases 'knowledge of French' and 'allergy to peanuts' denote states or properties rather than events, and so they cannot be combined with event-level adjectives. This explanation carries over straightforwardly to the adverbial case (25), on the assumption that 'knows French' and 'allergic to peanuts' are likewise not predicates of events. We can also make sense of the interaction between 'quickly' and negation observed in (27). On a standard event-based semantics, clausal negation always take wider scope than the event quantification, so that any event predicate occurring outside the scope of negation will lack an argument of the appropriate type; the interaction between 'quickly' and 'often' observed in (26) can be given an analogous explanation.[19] At the same time, it is no surprise that 'quickly' commutes with other constituents that also function semantically as event predicates, as observed in (29). Thus a great deal of the distinctive behavior of manner adverbs that is problematic for the operator account is easily explained on the hypothesis that they function as predicates of events.

However, we still need some account for non-intersectivity. Extant versions of the event analysis account for it by positing a second argument for comparison classes (Higginbotham 1985, Parsons 1990):

(36) a. Rolf swam quickly.
b. There was a swimming by Rolf that was quick for an *n*.

[18] (33b) also employs *thematic separation*, according to which the arguments of the verb – or more precisely, their occurrences in certain syntactic configurations – also function semantically as predicates of events. (For arguments see Parsons 1990 and Pietroski 2005.)

[19] One promising way to implement this is to interpret verbal projections, not as simple predicates of events, but as predicates of *sets* of events that apply just in case the set includes an event of the appropriate sort (Champollion 2015). This effectively treats existential quantification over events as already occurring within the verbal projection, so that cases like (26) and (27), in which 'quickly' occurs above the verbal projection, will be marked.

The value of '*n*' is taken to be fixed by context to determine a class of events with which Rolf's swimming is to be compared, such as a class of swimming events or a class of river crossings. But manner adverbs are also distinguished from modifiers like 'allegedly' in their ability to occur in comparative constructions and with degree modifiers like 'very' and 'somewhat':

(37) a. Rolf very/somewhat quickly (*allegedly) swam.
b. Rolf swam across the river more quickly (*allegedly) than Scooter did.

These phenomena interact with non-intersectivity. Notice, for example, that manner adverbs in comparative constructions no longer exhibit non-intersectivity: there is no sense to be made of the suggestion that Rolf swam more quickly than Gonzo for a swimming, but not for a river crossing.[20] This suggests that we should aim for a unified account of non-intersectivity, comparative constructions, and degree modification; unfortunately it is not obvious how such a treatment can be developed on the comparison class approach.

I suggest that we would do better by extending to manner adverbs a *degree-based* analysis of the sort already developed for gradable adjectives like 'large' (Cresswell 1976, Kennedy 1999, Kennedy and McNally 2005, von Stechow 1984). On this sort of approach, 'quickly' is taken to express a two-place relation between events and degrees, which are points on a scale that order events in terms of quickness, slowness, loudness, quietness and so on. The relation expressed by 'quickly' obtains between an event *e* and a degree *d* just in case the degree to which *e* is quick measures at least *d* on the scale. This allows us to interpret comparative constructions in terms of relations among degrees: (37b), for example, states that the maximal degree to which Rolf's swimming is quick is greater than every degree to which Scooter's swimming is quick. To say that someone swims quickly is to compare her swimming to some minimal threshold d_c set by context:

(38) a. Rolf swam quickly.
b. $\exists e$ (Swimming(e) & Agent(e,*rolf*) & $\exists d$ (Quick(e,d) & $d > d_c$)

Modifiers like 'very' and 'somewhat' can be understood as placing further conditions on the threshold d_c (or on the relation between d and d_c). On this approach, non-intersectivity is naturally seen as reflecting the role of context in selecting different values for the threshold: the value might be set one way in a context where we compare Rolf's swimming to others swimming events, and in another way in a context where we compare it to river crossings of all sorts.[21] Thus the degree-based analysis gives us a unified account of non-intersectivity and of comparatives and degree modification, while retaining the neo-Davidsonian's ability to explain the other kinds of behavior distinctive of manner adverbs.

Taken together, these observations provide us with a very strong case for regarding manner adverbs as a unified semantic natural kind, individuated by the property of expressing (denoting, standing for etc.) a relation between events and degrees; this is the underlying alethic essence that lexical items of this kind have in common.

It is straightforward to represent this within a Montagovian interpretational framework. Setting degrees aside for the moment, we can capture the idea that 'quickly' is an event predicate by first adding to our theory of types a new type **v** for (token) events. We then employ the standard Montagovian technique for representing predicative meanings via functions from predicate extensions to truth values. On the intended interpretation of 'quickly', for example, it is assigned a function of type <**v**,**t**> that takes an event as argument and returns *true* as its value if that event

[20] Non-intersectivity also interacts with 'very': an event might be quick both for a swimming and for a channel crossing, but only very quick for the former.
[21] Context might also play a role in selecting among different scales for measuring quickness.

is quick. On the neo-Davidsonian approach a verb like 'leave' also belongs to type **<v,t>**. This means that the semantic significance of the complex expression 'quickly leave' cannot simply be determined via function application, as on the operator approach. Instead we need to adopt a conjunctive composition rule to the effect that 'quickly leave' maps an event to *true* just in case the functions expressed by both 'quickly' *and* 'leave' map that event to *true*.

Things become slightly more complicated once degrees are brought back into the picture. We first add a new type **d** for degrees, and assign 'quickly' to the type **<d,<v,t>>**, i.e. functions from degrees to functions from events to truth values.[22] But this means that we no longer get a complex predicate of events simply by conjunctively combining the interpretations of 'quickly' and 'leave' as on the simpler picture; the degree argument of 'quickly' must first somehow be saturated. One way to do so is to hypothesize the presence of an unpronounced positive operator 'POS' that combines with 'quickly' to bind the degree argument and yield a predicate that applies to an event just in case the event is mapped onto some degree larger than a contextually-provided standard d_c. (This simply adapts a proposal from the adjectival case that traces back to Cresswell (1976) and von Stechow (1984).) This results in a predicate that can be conjunctively combined with 'leave' as before. Another way to do accomplish the same effect is to build the effect of 'POS' into the composition rule itself, so that combining 'quickly' with 'leave' simultaneously binds the degree argument of the former and conjoins it with the latter. There are no doubt other options as well.[23]

What is important for our purposes is that it is a more or less immediate consequence of adopting the conjunctive rule – in either its simpler or its more complex form – that 'Lisa leaves' will come out true (at a given context) relative to every admissible interpretation relative to which 'Lisa quickly leaves' is true (at that context). Thus on this account the factive entailment pattern comes out as categorially valid, and so genuinely structural. Since our hypothesis about the underlying alethic essence of manner adverbs is what leads us to adopt the conjunctive rule, we have here an illustration of the phenomenon noted in §3, where the effort to draw the correct distinctions among semantic kinds leads to adjustments in the composition rules, and this in turn leads naturally to the classification of some entailment patterns as structural.

What should we say about 'allegedly' and its cohort? It is not at all plausible that these share the alethic essence of manner adverbs: if they did then we would expect them to behave just like 'quickly' in all of the ways we have just seen they do not behave. In fact, it is much more plausible that 'allegedly' functions as a sentential operator that takes clausal complements, and whose semantic function is to express something like a property of propositions (Bellert 1977). On this sort of approach, 'Rolf allegedly left the stage' is roughly equivalent to 'It was alleged that Rolf left the stage'. This would explain all the behavior we have observed: that 'allegedly' is non-factive and intensional, that it is unrestricted in the types of predicates with which it occurs, that it exhibits scope interactions with negation and 'often', and that it cannot be used comparatively or with 'very'. Implementing this proposal within a Montagovian interpretational semantics is straightforward, and we will not undertake it here (see Thomason and Stalnaker (1973)). The important point for present purposes is that if this is correct then 'allegedly' has a very different

[22] The present proposal treats both events and degrees as basic types, rather than as higher types defined in terms of (for example) individuals and/or spatiotemporal regions and possible worlds. Eckardt (1998) presents what seems to me to be a strong abductive case for treating events as a basic type, and Penka and Balcerak Jackson (in preparation) makes observations about the explanatory role of degrees that encourage a treatment of degrees as a basic type. Needless to say, such treatments are compatible with various reductive metaphysical proposals about the nature of events and/or degrees.

[23] Conjunctive rules may turn out to have much wider use in semantics, in which case we may have good grounds for sticking with the simpler (and more generally applicable) rule and adopting the operator analysis instead (Pietroski 2005).

underlying alethic essence than 'quickly' – so different, in fact, that it is not easy even to see them as two sub-species of some higher-level kind to which we could give the label *adverb*.

Many questions remain open. Perhaps the most salient ones concern lexical items such as 'unexpectedly', 'intentionally', 'reluctantly', 'probably', and 'certainly' that do not clearly behave like either 'quickly' or 'allegedly'. Are these Lockean "freaks and monsters," deviant instances whose behavior diverges from the paradigm for the kind? Or do they constitute a further distinct natural kind, or perhaps several? How many natural kinds of modifiers are there in the verbal domain? On the present view these are ultimately questions about the underlying alethic essences of a large and diverse class expressions, and a comprehensive semantics of verbal modification would be much too ambitious of an undertaking for present purposes. The goal here has been simply to show that it is extremely plausible the natural kinds out of which 'Rolf quickly left the stage' and 'Rolf left the stage' is constructed, together with the role they play in determining truth conditions, conspire to guarantee that the former entails the latter. By present Evansian lights this makes it very plausible that the entailment is a genuinely structural one. More importantly, the discussion here has served to illustrate the way in which the present account gives substance to the question of whether a given entailment pattern counts as structural or non-structural. Even if a more comprehensive semantics for verbal modification ultimately leads us to classify the factive entailment as non-structural, this would do nothing to call into question the contention that the distinction itself is a genuine and significant one.

5 Epilogue

One might wonder: how important, ultimately, is the alleged *naturalness* of semantic natural kinds on the account developed here? Does it really require us to see semantic kinds as being very much like kinds in the natural sciences, as opposed to kinds of artifacts such as *chair* and *internal combustion engine*, or social or institutional kinds such as types of religious rituals or political systems? As stated, the question draws on a contrast between natural kinds and other categories that are somehow non-natural, and as far as I can see it matters little for the present account which side of this distinction semantic kinds are taken to fall on. In fact, I suspect that any attempt to draw a very sharp distinction here is likely to encounter trouble when applied to linguistic kinds. If semantic kinds are to play the role they are given here, what matters is that they be taken to be *objective*, in the sense that their existence and nature is not determined by the very activity of semantic inquiry itself, or by the decision to adopt a certain taxonomy for pragmatic reasons or as a matter of convention. Semantic kinds are not constructed by the theorist, but are "out there" for her to discover.

However, there is another, more methodological way to construe the question, and on this construal the conception of semantic kinds as natural kinds is much more important. The methodology employed in §4 for approaching questions about the taxonomy of adverbs is not so different than the chemist's methodology for approaching questions about kinds of inorganic compounds, or the biologist's methodology for approaching questions about morphological kinds. In each case kinds are discovered by finding diverse and *prima facie* unexpected patterns of shared behavior, and by identifying the underlying commonalities that best explain these patterns. This is ultimately why it is misguided to enforce a general presumption in favor of uniform analyses in semantics. Preserving such analyses inevitably forces one to dismiss differences in entailments or other semantically relevant behavior as mere lexical variations, and this frustrates the discovery of the genuine joints in semantic reality that is, on the present view, one of the central aims of semantic theorizing. Of course, simplicity and parsimony play a role in semantic theory choice in semantics, as elsewhere. But the ability to offer parsimonious *unifying* explanations of a wide and diverse range of phenomena is a much more important theoretical

virtue than the ability to offer uniform explanations in terms of a small number of coarsely individuated kinds.

This brings us back, finally, to the source of the disappointing results in Montague's EFL. For Montague's strategy of generalizing to the worst case simply *is* a version of the strategy of pursuing uniform explanations at the expense of unifying explanations. The fact that the semantics in EFL fails to capture any interesting structural entailment patterns is the predictable result of applying this strategy. In other words, what is missing in EFL is the recognition that the basic categories out of which a semantic theory is to be constructed should be understood as kinds that are to be investigated in essentially the same way as kinds in chemistry, biology and other natural sciences. EFL is a beautiful demonstration that the tools of the logician can be applied to natural language. But to be fruitful they need to be wielded according to the methods of the natural scientist.

References

Balcerak Jackson, Brendan. 2007. Beyond logical form. *Philosophical Studies* 132, 347-380.

----. 2006. Logical form: classical conception and recent challenges. *Philosophy Compass* 1, 303-316.

Barwise, Jon and Robin Cooper. 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy*, 4, 159-219.

Bealer, George. 1982. *Quality and Concept*. Oxford: Clarendon Press.

Bellert, I. 1977. On semantic and distributional properties of sentential adverbs. *Linguistic Inquiry*, 8, 337-351.

Boyd, Richard. 1999. Homeostasis, species, and higher taxa. In R.A. Wilson (ed.), *Species: New Interdisciplinary Essays*. Cambridge: MIT Press.

----. 1989. What realism implies and what it does not. *Dialectica*, 43, 5-29.

Chalmers, David J. 2002. On sense and intension. *Philosophical Perspectives*, 16, 135-182.

Champollion, Lucas. 2015. The interaction of compositional semantics and event semantics. *Linguistics and Philosophy* 38, 31-66.

Cresswell, M.J. 1976. The semantics of degree. In B. Partee (ed.), *Montague Grammar*. New York: Academic Press.

----. 1974. Adverbs and events. *Synthese*, 28, 455-481.

Davidson, Donald. 1973. In defence of Convention T. In H. Leblanc (ed.), *Truth, Syntax and Modality*. Dordrecht: North-Holland Publishing Company.

----. 1970. Semantics for natural languages. In C. Olivetti (ed.), *Linguaggi nella Societa e nella Tecnica*. Milan: Edizioni di Comunità.

----. 1967. Truth and meaning. *Synthese*, 17, 304-323.

----. 1966. Theories of meaning and learnable languages. Y. Bar-Hillel (ed.), *Proceedings of the 1964 International Congress for Logic, Methodology and Philosophy of Science*. Amsterdam: North Holland Publishing Company.

----. 1967. The logical form of action sentences. In N. Rescher (ed.), *The Logic of Decision and Action*. Pittsburgh: University of Pittsburgh Press.

Dummett, Michael. 1991. *The Logical Basis of Metaphysics*. Cambridge: Harvard University Press.

Eckardt, Regina. 1998. *Adverbs, Events, and Other Things: Issues in the Semantics of Manner Adverbs*. Tübingen: Max Niemeyer Verlag.

Evans, Gareth. 1985. Semantic structure and logical form. In *Collected Papers*. Oxford: Oxford University Press.

Higginbotham, James. 1992. Truth and understanding. *Philosophical Studies*, 65, 3-16.

----. 1985. On semantics. *Linguistic Inquiry*, 16, 547-593.

Horwich, Paul. 2001. Deflating compositionality. *Ratio*, 14, 369-385.

Jackendoff, Ray S. 1990. *Semantic Structures*. Cambridge: MIT Press.

Johnson, Kent. 2003. Are there semantic natural kinds of words? *Mind & Language*, 18, 175-193.

Katz, Gareth. 2003. Event arguments, adverb selection, and the Stative Adverb Gap. In E. Lang, C. Maienborn & C. Fabricius-Hansen (eds.), *Modifying Adjuncts*. Berlin: Mouton de Gruyter.

Kempson, Ruth. 2011. Formal semantics and representationalism. In C. Maienborn, K. von Heusinger & P. Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning*. Berlin: Mouton de Gruyter.

Kennedy, Christopher. 1999. *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. New York: Garland Publishing, Inc.

Kennedy, Christopher and Louise McNally. 2005. Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81, 345-381.

King, Jeffrey C. 2007. *The Nature and Structure of Content*. Oxford: Oxford University Press.

Levin, Beth. 1993. *English Verb Classes and Alternations*. Chicago: The University of Chicago Press.

Lewis, David. 1975. Languages and language. In K. Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*. Minneapolis: University of Minnesota Press.

Lidz, Jeff, Paul Pietroski, Tim Hunter and Justin Halberda. 2011. Interface transparency and the psychosemantics of 'most'. *Natural Language Semantics* 19, 227-256.

Ludlow, Peter. 2002. LF and natural logic. In G. Preyer & G. Peter (eds.), *Logical Form and Language*. Oxford: Oxford University Press.

McConnell-Ginet, Sally. 1982. Adverbs and logical form: a linguistically realistic theory. *Language*, 58, 144-184.

McGilvray, James. 1998. Meanings are syntactically individuated and found in the head. *Mind & Language*, 13, 225-280.

Montague, Richard. 1974. English as a formal language. In *Formal Philosophy: Selected Papers of Richard Montague*. New Haven: Yale University Press.

Parsons, Terence. 1990. *Events in the Semantics of English: A Study in Subatomic Semantics*. Cambridge: MIT Press.

Penka, Doris and Brendan Balcerak Jackson. In preparation. Number sentences and the metaphysics of degrees.

Pietroski, Paul. 2005. *Events and Semantic Architecture*. New York: Oxford University Press.

Pietroski, Paul, Jeff Lidz, Justin Halberda and Tim Hunter. 2009. The meaning of 'most': semantics, numerosity, and psychology. *Mind and Language* 24, 554-585.

Salmon, Nathan. 1986. *Frege's Puzzle*. Cambridge: MIT Press.

Soames, Scott. 1987. Direct reference, propositional attitudes and semantic content. *Philosophical Topics*, 15, 47-87.

von Stechow, Arnim. 1984. Comparing semantic theories of comparison. *Journal of Semantics*, 3, 1-77.

Szabó, Zoltán Gendler. In preparation. Semantic explanations.

Thomason, Richmond and Robert Stalnaker. 1973. A semantic theory of adverbs. *Linguistic Inquiry*, 4, 195-220.

Williams, Edwin. 1983. Semantic vs. syntactic categories. *Linguistics and Philosophy*, 6, 423-446.

Wright, Crispin. 1981. Rule-following, objectivity and the theory of meaning. In S. H. Holtzman & C. M. Leich (eds.), *Wittgenstein: To Follow a Rule*. Oxford: Routledge.

Zimmermann, Thomas Ede. 1999. Meaning postulates and the model-theoretic approach to natural language semantics. *Linguistics and Philosophy*, 22, 529-561.