

Adjectives and Headedness*

Robert Truswell

1 Introduction

Within the Principles and Parameters framework, two standard assumptions concerning the syntactic and semantic nature of attributive adjectives are that they denote properties, and that they occupy adjoined positions.¹ The former assumption found a natural home in strongly first-order theories such as Higginbotham 1985 and Chierchia & Turner 1988, although it also formed a central instance of the operation of Predicate Modification in Heim & Kratzer 1998. Meanwhile, the latter assumption appeared to go without saying in early Government and Binding theory: attributive adjectives, as optional elements, would appear to be prime candidates to occupy adjoined positions.

This paper will argue, however, that in the case of English, both of these assumptions are incorrect. It will be demonstrated that the major alternative semantic treatment of attributive adjectives, as denoting second-order identity-typed functions, has conceptual and empirical advantages over a theory in which attributive adjectives denote properties. If, moreover, the requirement that an attributive adjective modify a nominal is a specifically syntactic selectional requirement, this would suggest, following Chomsky 2000 among others, that it is the adjective, rather than the noun, which projects in this construction. In that case, the adjective has one of the major characteristics of syntactic heads, and so the analysis of attributive adjectives as heads in Abney 1987 gains plausibility.

The following sections will spell out the steps in the above argument in greater detail. §2 will compare the analyses of attributive adjectives as first- and second-order functions, and the predictions that they make. §3 will argue that the need for an attributive adjective to modify a noun does not stem from the compositional

*Thanks to Ad Neeleman for comments on an earlier draft. This paper has its roots in my M.Phil thesis. I am grateful to Gillian Ramchand, the supervisor of that thesis, and to Peter Svenonius, for helpful suggestions and encouragement. Section 4.1 is based on work by, and discussion with, Michael Wagner. I am very grateful to him for sharing his work and ideas with me. All remaining errors are my own.

¹This also holds in work, stemming from the Linear Correspondence Axiom (Kayne 1994) and developed by Cinque 1994, which takes attributive adjectives to occupy specifier positions within an extended nominal functional sequence, as the LCA characterises specifier positions as the unique adjoined maximal projections within each phrase.

semantics, but rather that there is a parallel syntactic selectional requirement for attributive adjectives to have a nominal sister. This section will also sketch one way of accommodating such a requirement within current syntactic theory. Finally, §4 will show some further empirical advantages of this syntactic treatment.

2 Two Semantic Analyses of Attributive Adjectives

2.1 The Montagovian Analysis

Within the framework of Montague Grammar, the assumption that attributive adjectives are second order functions is absolutely natural. Distributionally, an Adj–N group is identical to a common noun in isolation. If the syntactic category of common nouns is CN, then, the natural analysis of attributive adjectives is as being CN/CN elements. On the assumption that common nouns are of type $\langle e, t \rangle$, then, attributive adjectives are of type $\langle \langle e, t \rangle, \langle e, t \rangle \rangle$, their identity type reflecting their syntactic transparency. Indeed, *modulo* intensionality, this is exactly the proposal of Lewis 1972, Montague 1974 and Kamp 1975.

Kamp defines several subclasses of attributive adjective on the basis of the relations that hold between $\llbracket \text{Adj} \rrbracket$, $\llbracket \text{N} \rrbracket$ and $\llbracket \text{Adj} \rrbracket(\llbracket \text{N} \rrbracket)$. For example, an adjective F is *intersective*² if:

- (1) ‘there is a property Q such that for each property P and each $w \in W$ [the set of possible worlds], $F(P)(w) = P(w) \cap Q(w)$ ’ (Kamp 1975:124);

and an adjective is *affirmative* if:

- (2) ‘For each P and w, $F(P)(w) \subseteq P(w)$ ’ (Kamp 1975:125).

Crucially, though, these properties of classes of adjectives are irrelevant to their combinatorial properties: they are meaning postulates, or more specifically, statements about our knowledge of the relation between elements in the domain and the range of a given adjectival function.

2.2 On Semantics and All that Followed


The Montagovian treatment of attributive adjectives posits a uniform method of syntactic and semantic combination whereby adjectives are functions from property-denoting common nouns into property-denoting common nouns. The different entailment relations that different attributive adjectives give rise to must then be treated as meaning postulates, as in (1–2), concerning the relation between the domain and range of a given adjective.

²I substitute the more current term *intersective* for Kamp’s *predicative*, which is used in a distinct sense below.


An alternative view was proposed by Higginbotham 1985, arguing for an explicit representation of these relations in the syntactic and semantic derivation. Higginbotham assumes that an attributive adjective, as well as a noun and an Adj-N group, denotes a property. Since functor-argument relations between two $\langle e, t \rangle$ elements as they stand are impossible, as neither element is of the correct type to serve as argument to the other, the immediate question is one of how the adjective and noun are to combine semantically.

Higginbotham's solution is to shift the relationship between adjective and noun into θ -theory. He assumes that his proposal that adjectives and nouns denote first-order functions translates into such elements bearing θ -roles, which must be discharged. There are two mechanisms by which discharge of adjectival θ -roles is achieved, namely θ -identification and *autonomous θ -marking*. These correspond, roughly, to Kamp's properties of *affirmativity* and *non-intersectivity* respectively, but here, they are structurally encoded. Diagrammatically, these operations are represented by Higginbotham as follows:

(3) a. θ -identification:



b. Autonomous θ -marking:



We can then define three classes of adjective, according to the thematic relations into which a given adjective enters:

- (4) a. An adjective whose θ -role is discharged through θ -identification is affirmative and intersective.
- b. An adjective whose θ -role is discharged through autonomous θ -marking is non-affirmative and non-intersective.
- c. An adjective with two θ -roles, one discharged through θ -identification and one through autonomous θ -marking, is affirmative and non-intersective.

The fourth logical possibility, an adjective which is non-affirmative and intersective, can never arise, as there is a contradiction between these two characteristics: as defined above, intersectivity entails affirmativity. Also, in θ -theoretic terms,

such an adjective would not have entered into a relation of θ -identification or autonomous θ -marking, and so it would have an undischarged θ -role, in violation of the θ -criterion.

Broadly, then, we may suppose that the semantic representations of these three types of adjective–noun groups, respectively, will be as follows:

- (5) a. Affirmative and intersective: $\lambda x.N(x) \wedge A(x)$
- b. Non-affirmative and non-intersective: $\lambda x.A(\wedge N(x))$ ³
- c. Affirmative and non-intersective: $\lambda x.N(x) \wedge A(\wedge N(x))$

2.3 Separating the Two Analyses

The theory represented in (4–5) predicts that inferences parallel to the meaning postulates in (1) and (2) are structurally encoded in the syntactic and semantic derivation. If, for example, an adjective belongs to the class (5a), the inferences for some individual x from $\text{Adj-N}(x)$ ⁴ to $\text{Adj}(x)$ and $N(x)$ should then be automatic. Similarly, if an adjective is in class (5c), the inference from $\text{Adj-N}(x)$ to $N(x)$ should be automatic. Any other inferences, for adjectives in one of the classes in (5), should be invalid. Such patterns of inference clearly do not necessarily hold so absolutely under a view of affirmativity and intersectivity as meaning postulates.

Indeed, this is part of the motivation for Higginbotham’s move away from the Montagovian view of attributive adjectives sketched in §2.1: if we can automatically infer that a *red car* is both *red* and a *car*, and that a *big car* is a *car*, but not that a *big car* is *big* or a *possible car* is a *car*, shouldn’t the automatic inferences be “hard-wired” into the compositional semantics?

The question, then, is whether these “automatic” inferences really are automatic. And it would seem that they are not. Consider, first, the following two phrases:

- (6) a. A red face
- b. A red double-decker bus

Clearly, the *red* in these two examples is not identical, as a *red face* and a *red double-decker bus* are probably not the same colour. This means that even a prototypically intersective adjective, such as *red*, does not simply denote a property in (6), but is relativised to the noun that it modifies. In other words, even adjectives

³Higginbotham’s proposal is that what is θ -marked in autonomous θ -marking is ‘the phrase marker with root N ’ (Higginbotham 1985:564), as opposed to the reference of such a phrase marker. It is not clear to me what this corresponds to in the model, and so I represent autonomous θ -marking here as taking the nominal intension as an argument.

⁴This notation is intended to be neutral between the two representations, $(\text{Adj}(N))(x)$ and $\text{Adj}(x) \wedge N(x)$.

such as *red* are not absolutely intersective, in the sense of (1). This casts some doubt on the existence of (5a) as a separate class of adjectives.

It may be argued, though, again following Higginbotham 1985, that a *red face* is still a *face*, and a *red double-decker bus* is still a *double-decker bus*. In the above examples, this is undoubtedly true. However, consider the following sentences, all of which I would claim are true:

- (7) a. A corn marigold is not an ox-eye daisy. [They have separate Latin names, for example]
- b. A corn marigold is (just) a yellow ox-eye daisy.⁵
- c. Therefore, a yellow ox-eye daisy is not (necessarily) an ox-eye daisy.
- (8) a. A marrow is not a courgette. [They are sold as separate vegetables by greengrocers, for example]
- b. A marrow is (just) a big courgette.
- c. Therefore, a big courgette is not (necessarily) a courgette.
- (9) a. A Tonka truck is not a truck. [This will become evident if you attempt to drive one on a public highway]
- b. A Tonka truck is a toy truck.
- c. Therefore, a toy truck is not (necessarily) a truck.

In each case, the pattern is the same. The (a) sentences assert that the two noun phrases denote distinct properties, while the (b) sentences assert that one particular property, denoted by the adjective, distinguishes the properties denoted by the two noun phrases⁶. This allows us to deduce, as in the (c) sentences, that the inference from Adj-N to N is not automatic.

However, only (9) involves an adjective which would standardly be considered as non-affirmative, where *affirmative* is defined as in (2). (7–8) involve canonical intersective and subsective adjectives, respectively.⁷ Under a theory of attributive adjectives such as Higginbotham's, the failure of the inference that "an Adj N is an N" is unexpected in such cases, as such inferences are inseparably linked, through

⁵I have no explanation for why it should be that *just* facilitates interpretation of the (b) examples.

⁶There is a sense in which the (a) and (b) sentences are true at different levels of granularity. It is only at a fairly superficial level that the equation "ox-eye daisy + yellow = corn marigold" could be taken to be true. That is beside the point, here, touching on questions of how an audience evaluates the truth of such utterances. The fact that they *could* all be taken to be true, under quite unremarkable circumstances, already has serious implications for the notion that attributive adjectives can denote conjoined properties.

⁷A *subsective* adjective is here defined as a non-intersective, affirmative adjective, that is, the class (5c).

θ -theory, to the fundamental mode of structural combination of such adjectives with nouns. However, if the combination of adjective and noun is independent of such entailment relations, which are instead represented as meaning postulates, as in the Montagovian theory, then their defeasibility in contexts such as (7–9) has less drastic theoretical consequences. Such examples must be taken, then, to strongly favour the uniform analysis of attributive adjectives as second-order functions, and to provide evidence against θ -identification as the mode of combination of attributive adjectives.

2.4 Relating Attributive and Predicative Uses of Adjectives

A consequence of Higginbotham's claim that adjectives canonically denote properties is that we might then expect a primary function of theirs to be forming propositions by taking type e arguments. Indeed, the evidence from copular and predicative constructions suggests that we should preserve this intuition, as Partee 1987 demonstrates that adjectives in such constructions alternate with indefinite DPs, and can be co-ordinated with them:

- (10) Mary considers John competent in semantics and an authority on unicorns.
(Partee 1987:119)

On the other hand, the evidence reviewed above from attributive uses suggests that adjectives can be second-order functions, of type $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$. If this is the case, we may assume that an operator is available of type $\langle\langle e, t \rangle, \langle\langle e, t \rangle, \langle e, t \rangle\rangle\rangle$, that is, an operator which takes a first-order property as its argument and outputs a second-order function from properties to properties, shifting from the predicative to attributive use of an adjective.

Postulation of such an operator permits explanation of several facts concerning the distribution of adjectives. For example, it allows us to suggest a semantic reformulation of the distinction proposed in Bernstein 1993 and Alexiadou 2001 between a class of adjectives that are X^0 s and a class that are XPs: in Bernstein's analysis, the A^0 s are distinguished by never occurring in predicative constructions and by not taking modifiers such as *very*, for example:

- (11) a. 'That car is big'
b. 'A very big car'
c. * 'That car is former'
d. * 'A very former car'

Such facts are amenable to a more structurally uniform formulation if we assume that, in fact, those adjectives which Bernstein labels as A^0 s are lexically represented as type $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$, while those that she labels as APs are represented

as type $\langle e, t \rangle$ and require a type-shifting operator to be used attributively. If, furthermore, degree modifiers take $\langle e, t \rangle$ arguments (which seems reasonable, given the acceptability of a sentence such as ‘That car is [very big]’), we have largely derived the restrictions on adjective distribution in (11).

Furthermore, postulation of this operator goes some way towards resolving the tension observed in Kamp 1975 between the fact that attributive uses of adjectives are most satisfactorily represented as second-order functions, while a supervaluation-based treatment of comparatives seems to require that adjectives denote gradable first-order properties. As the cases of comparatives he discusses are all found in clausal constructions, rather than DP-internally, this distinction is now expected. We see that it is quite generally the case that predicative uses of adjectives are first order, and attributive uses are second order, and the comparative constructions that Kamp discusses fit this pattern.

As an approximation of the content of this operator, I propose to reformulate the JOIN operator of Baker 2003 in such a way as to allow a second-order representation of attributive adjectives:


(12) $\text{JOIN}(A) = \lambda N \lambda x. (A_{\text{Attr.}}(\wedge N))(x)$, where:

A is the property denoted by the adjective used predicatively;

$A_{\text{Attr.}}$ is the second-order function denoted by the adjective used attributively;

N is the property denoted by the noun (with any lower complements or modifiers).⁸

The syntax and semantics of attributive adjectival modification under these proposals, putting aside for the moment issues of projection, will be as follows:



b. $(\text{JOIN}(\text{red}'))(\text{car}') = \lambda y. (\text{red}_{\text{Attr.}}(\text{car}'))(y)$

This representation captures several essential facts about attributive adjectives: the modification is potentially recursive because the type of the noun is unchanged by merging an adjective; the modification is optional because what is usually essential for a DP is that it provides an argument to the clause which contains it, and this is neither helped nor hindered by merging an adjective because there is nothing in the nominal semantics which requires the merging of an adjective (the requirement comes instead from JOIN, which requires two $\langle e, t \rangle$ elements as arguments);

⁸This formulation clearly begs the question of the nature of the almost, but not quite, deterministic relation between an adjectival property and the related second-order function. This important problem will have to await further research.

and scope is represented because the most recently merged adjective directly modifies the group consisting of the noun and any adjectives merged earlier, and this group is basically semantically opaque (i.e. an Adj–N group behaves exactly like a bare noun as far as any further computation is concerned).

It is important to note that postulation of JOIN means that, even in cases such as English where the phonetic forms of adjectives used predicatively and attributively are identical, they make distinct contributions to the semantic representation. In particular, a predicative adjective is first-order and an attributive adjective is second-order. Note also that JOIN must be formulated in such a way as to take the adjective as its first argument, in order to capture the data in (11). A possible alternative formulation of JOIN which took the adjectival and nominal arguments in the opposite order would necessarily assume a uniform type for adjectives, ruling out the possibility that certain adjectives could be lexically specified as type $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$. These considerations will have repercussions when we turn our attention to syntax in the following section.

3 English Attributive Adjectives are Heads

We have seen that attributive adjectives denote functions from properties to properties, and that they are derived from first-order properties by means of the two-place relation JOIN. However, not every property constitutes an acceptable input to such a relation. For instance, we also saw in §2.4 above that there is reason to believe that predicative adjectives denote properties, yet an attributive adjective cannot modify a predicative adjective:

(14) a. *The car was *big red*.

$$\text{b. } (\text{JOIN}(\text{big}'))(\text{red}') = \lambda y.(\text{big}_{\text{Attr.}}(\text{red}'))(y)$$

Furthermore, JOIN must only be able to take an adjectival property as its first argument, otherwise the following should be derivable, by JOIN taking two nominal arguments:

(15) a. *The *car vehicle* raced down the street.

$$\text{b. } (\text{JOIN}(\text{car}'))(\text{vehicle}') = \lambda y.(\text{car}_{\text{Attr.}}(\text{vehicle}'))(y)$$

As the only elements modified by attributive adjectives are nominals (that is, a noun with any complements, possibly modified by other, lower, adjuncts), and as an attributive adjective with no nominal sister is marked, if not ungrammatical (compare null nominal constructions such as (16)), it appears that an attributive adjective selects an N projection as its sister.⁹


⁹Note that such a suggestion is only plausible if attributive and predicative uses of adjectives are syntactically distinct. Predicative adjectives do not subcategorise for a sister at all, instead standardly occurring as the sister of a predicative head.

(16) # The red is good.

Now, under a standard conception of attributive adjectives as adjuncts, the nominal projects, by definition. For instance, Chomsky 2000 writes that:

‘Adjunction has an inherent asymmetry: X is adjoined to Y. Exploiting that property, let us take the distinction between substitution and adjunction to be the (minimal) distinction between the set $\{\alpha, \beta\}$ and the ordered pair $\langle \alpha, \beta \rangle$, α adjoined to β . . . [In the latter case,] Given the asymmetry, it is natural to conclude that the adjoined element α leaves the category type unchanged: the target β projects.’ (Chomsky 2000:133)

Leaving aside the question of whether JOIN has a syntactic reality, or instead perhaps describes a lexical redundancy rule, this would give a tree for an Adj–N constituent as follows:




With regard to projection in cases of substitution, Chomsky writes:

‘Set-Merge typically has an inherent asymmetry. When α, β merge, it is to satisfy (selectional) requirements of one (the *selector*) but not both. Fairly generally, furthermore, the selector is uniquely determined for a pair (α, β) . . . In this case too, then. . . the label of the selector projects.’ (Chomsky 2000:133–4)

These two statements, however, should give rise to doubts concerning the validity of the substitution–adjunction distinction, formulated in this way. It is the inherent asymmetry in adjunction which leads Chomsky to propose that adjunction is a pair-forming operation. However, there is an inherent asymmetry in substitution too. It is not clear why these two asymmetries should be represented in different ways (ordering of elements and projection in the case of adjunction, as opposed to projection alone in the case of substitution). Even if it is indeed necessary to preserve the substitution–adjunction distinction (and much recent research, from Kayne 1994 onwards, suggests that it is not), it is at least clear that the way of representing this distinction in Chomsky 2000 amounts to nothing more than a diacritic, and so should be treated with suspicion, in line with minimalist principles.

Let us assume, instead, that any ‘inherent asymmetry’ is represented structurally by projection alone. If adjunction is not treated as a structurally distinct operation from substitution, then a surprising reversal occurs. It is clear that the nominal is Merged to satisfy the selectional requirements of the attributive adjective, rather than *vice versa*: it is the attributive adjective which subcategorises for a nominal sister, while the nominal clearly does not subcategorise for an adjective. In that case, we should expect the adjective, rather than the noun, to project:



In other words, if it is possible to equate selection with projection, the head of an Adj–N constituent is the adjective. However, the distribution of Adj–N constituents is identical to that of nouns alone. It is necessary, then, to explain how the adjective projection comes to be essentially “transparent”, in that it allows categorial features of its nominal sister to remain visible.

3.1 Attributive Adjectives and Relativised Heads

Baker 2003 proposes a theory in which there are three lexical categories, namely nouns, verbs and adjectives / adverbs. Nouns are distinguished by having a *criterion of identity* (that is, ‘they... set standards by which one can judge whether two things are the same or not’, Baker 2003:101), and verbs by their ability to license specifiers. Adjectives and adverbs are distinguished by the absence of both of these characteristics. The claim, then, is that it is only meaningful to say *X is the same Y as Z* or *X is a different Y to Z* if Y has the distinguishing characteristic of a nominal. If an Adj–N constituent is to have the function and syntactic distribution of a noun, then, it follows that the Adj–N constituent, too, must be able to set a standard of sameness. This appears to be true:

(19) Cabbage is a different green vegetable to kale.

If an Adj–N constituent, like a noun, has a criterion of identity, and this is the fundamentally nominal characteristic, then the fact that attributive adjectives do not alter the distribution and function of their nominal sister is unsurprising. This criterion of identity must, however, be inherited from the noun, rather than the adjective, as adjectives are not inherently able to express set standards of sameness. This is surprising, given the evidence that the adjective is the head of an Adj–N constituent.

This phenomenon, whereby syntactic characteristics are inherited from the non-head, is described extensively in Williams 1994, where the non-head is called a *relativised head*. Williams’ proposal is as follows:

‘Suppose first that there is an “absolute” head. Then, “head with respect to F” is defined as follows:

(45) X is the head with respect to F of Y if X is marked for a value of F, and either X is the absolute head of Y, or the absolute head of Y is not marked for F.

This provides a sort of unification that always succeeds, for in cases of conflict the head wins.’ (Williams 1994:46)

Assume that the criterion of identity is encoded by some feature. As this is essentially a variant of the long tradition of describing lexical categories in terms of categorial features, let us call the feature [N]. Then nouns have this feature, and adjectives don't, as a direct consequence of Baker's theory. Furthermore, the [N] feature on a nominal complement can only be inherited by a head if that head is not verbal, as no category can be both nominal and verbal on Baker's theory. In that case, in the construction [_{AP} A [_{NP} N]], A can inherit N's [N] feature, as A is defined as a lexical category lacking the defining characteristics of both nouns and verbs. This contrasts with cases where, for example, a verb has a nominal complement. As [V] and [N] features are incompatible, there is a conflict, and the absolute head alone determines the category of [_{VP} V DP].

We must note, however, that it is possible to subcategorise for an adjectival projection, as in the case of *too*, for example, which only occurs modifying adjectival and adverbial projections. This means that an adjective cannot simply be analysed as a lexical category without [N] and [V] features, as, in that case, there will be nothing for *too* to subcategorise for. However, there is evidence that it is only possible to subcategorise for *predicative* adjectives in this way: the sequence D-*too*-A-N is at least marginal in English, and adjectives argued above to be lexically specified as type $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$ are unable to be modified by *too*:

- (20) a. ?? A too happy man
 b. * Too former / too solar

This constitutes further evidence that the syntactic nature of predicative and attributive adjectives is radically different: to the best of my knowledge, there is no clear evidence that anything subcategorises for a specifically attributive adjective in English. The categorial status of an attributive adjective with respect to the noun that it modifies is reminiscent, then, of the status of functional heads: although they head their own projections, these still form part of the *extended projection* (in the sense of Grimshaw 2003) of the lexical item most locally c-commanded by the functional head. This paper claims that attributive adjectives are heads in the nominal extended projection, even if predicative adjectives are associated with some distinct categorial feature, which permits elements such as *too* and *very* to subcategorise for them.

The theory of projection and category membership sketched here amounts to saying that, no matter which daughter projects, if the mother can function as a criterion of identity, then it will have the distribution of a noun phrase. And only nominal extended projections (by definition) and projections of attributive adjectives (where the ability to function as an expression of sortality is inherited from the nominal complement, there being no conflict between an attributive adjective and this nominal characteristic) can function in this way.

At least in English, one welcome immediate consequence of such an analysis is that the general Adj-N order then coincides with the general head-complement

order. However, DP-internal adjectives are not always pre-nominal in English: notably, adjectives with PP complements occur post-nominally, and the same applies to occurrences within the noun phrase of adjectives modified by *too*:

- (21) a. A proud man
b. * A man proud
c. * A proud of his children man
d. A man proud of his children

- (22) a. * A too proud man
b. A man too proud

Interestingly, Williams himself assumes that certain adjectives, notably *alleged*, stand in a head-complement relation to the nouns they modify. One reason for singling out *alleged* for such an analysis is that it has distinct semantic characteristics when in post-nominal position. While an *alleged murderer* is not necessarily a murderer, an instance of *murderer* modified by post-nominal *alleged* necessarily does denote a murderer:

- (23) a. * The murderer alleged since yesterday. . .
b. The murderer alleged to have stolen the car. . . (Williams 1994:92)

However, given the discussion in §2.3, we may now expect to find that such patterns hold more generally. And it seems that this is true: the following examples must denote ox-eye daisies and courgettes, respectively, in contrast to the pre-nominal cases:

- (24) a. (# A corn marigold is just. . .) an ox-eye daisy, yellow as the sun.
b. (# A marrow is just. . .) a courgette as big as a shoe.

One interesting adjective in this respect is *old*. *Old* is vague: it can prototypically be interpreted affirmatively (in which case it is the opposite of *young*), or non-affirmatively (in which case its meaning is close to *former*)¹⁰. However, used post-nominally, only the first sense remains: the following must necessarily denote a house:

- (25) A house as old as the hills. . .

¹⁰In fact, the discussion in §2.3 leads us to expect that non-affirmative uses with the former sense should also be possible, and this prediction is, indeed, borne out, e.g.: *Stilton is just old cheddar*.

This evidence all suggests that post-nominal adjectives display all the characteristics of predicative, not attributive, semantics: N–Adj constructions, strikingly, allow affirmative patterns of inference and permit the use of modifiers available in predicative constructions but barred pre-nominally. We may speculate, then, that the N–Adj construction is structurally distinct from the Adj–N construction, in that only the latter makes use of JOIN.

In that case, generalising Williams’ analysis of *alleged* to cover all attributive adjectives is theoretically plausible, and, it seems, empirically justified. A pre-nominal attributive adjective is a head, and the noun is its complement. In the following section, I will sketch some further benefits of this approach.

4 Extensions

4.1 Adjectives, Givenness and Prosody

Givenness of material is generally marked in English by de-accenting (roughly in the sense of absence of a pitch accent where one may otherwise be expected), as can be seen by contrasting the following two exchanges, where capitals denote a pitch-accented syllable:

- (26) After Mary came, what happened next?
- a. She PRAISED JOHN.
 - b. # She PRAISED John. (Wagner 2005:3)
- (27) Mary met John. What happened then?
- a. # She PRAISED JOHN.
 - b. She PRAISED John. (Wagner 2005:6)

In (26), *John* is new material, and so de-accenting it in the response (26b) is odd. On the other hand, in (27), *John*, and more specifically the event of Mary meeting John, is salient, having featured in the question. This means that a structure in which an event of the form *Mary X John* is presented as given, and only the fact that it is a praising event is presented as new, as in (27b), is possible.

In theories of givenness since Schwarzschild 1999, givenness has been defined in terms related to Selkirk’s theory of intonational focus. That theory assumes that pitch accent on a word *F-marks* that word, and that F-marking of units containing that word is governed by rules of *focus projection*. The focus of a sentence is then identified with the maximal F-marked unit within that sentence:

- (28) a. F-marking of the *head* of a phrase licenses the F-marking of the phrase.

- b. F-marking of an *internal argument* or a head licenses the F-marking of the head. (Selkirk 1995:555)

Crucially, in this theory, F-marking may only spread to a head X from complements of X, not from specifiers or adjuncts. This derives the contrast between the different foci corresponding to F-marking on the complement and focus projection as governed by (28) (shown below by different acceptable congruent questions), and the contribution of F-marking on the subject being restricted to subject focus:

(29) [Mary [bought [a book [about [BATS]_F]]]]

- a. What did Mary buy a book about?
- b. What kind of book did Mary buy?
- c. What did Mary buy?
- d. What did Mary do?
- e. What's been happening? (Selkirk 1995:554)

(30) [MARY]_F bought a book about bats

Who bought a book about bats? (Selkirk 1995:554)

Note, firstly, that attributive adjectives pattern like heads, rather than specifiers or adjuncts, in this respect:

(31) [Mary bought [a big [red [CAR]_F]]]

- a. What big red thing did Mary buy?
- b. What big thing did Mary buy?
- c. What did Mary buy? (etc. . .)

Equally, attributive adjectives pattern like heads with respect to the requirement on deaccenting of given material:

(32) A: Why don't you have some French TOAST?

B: I've forgotten how to MAKE French toast (Schwarzschild 1999:142)

(33) {John drove Mary's red convertible. What did he drive before that?}

A: He drove her [BLUE]_F convertible. (Schwarzschild 1999:146)

It seems that the patterns of F-marking, focus and givenness holding between an adjective and a nominal exactly parallel those holding between a head and a complement. Of course, this comes for free if the nominal *is* the complement of the phrase containing the adjective.¹¹ Adopting an analysis in which the attributive adjective phrase is a head, then, allows us to capture such information-structural effects within the noun phrase in a natural way.

4.2 Adjectives and Head Movement

Consider the following data from Danish. Danish definite articles have two forms, a suffixal form occurring on nouns unmodified by attributive adjectives, and a full word which appears when the noun phrase contains one or more adjectives:

- (34) a. Hus -et
 House -the
 'The house'
- b. Det lange hus
 'The long house'
- c. *(Lange) huset (lange) (data adapted from Vangsnes 1999:102)

An appealing analysis of this alternation relies on N-to-D raising to generate (34a). If *-et* requires an N^0 host to attach to, then an attributive adjective phrase, if analysed as the head of the Adj-N constituent, should count as an intervener for the Head Movement Constraint, hence the ungrammaticality of (34c). This could, in turn, require use of the stand-alone article *det*, as in (34b):


- (35) a.
-
- ```

graph TD
 DP --> D
 DP --> NP
 D --> Hus_i
 D --> et["-et"]
 NP --> t_i

```
- b. \*
- 
- ```

graph TD
  DP --> D
  DP --> AdjP
  D --> Hus_i
  D --> et["-et"]
  AdjP --> lange
  AdjP --> NP
  NP --> t_i
  
```

¹¹Schwarzschild offers an alternative analysis, preserving the adjunction analysis of adjectives, and relying on entailment relations to derive these patterns. While space prevents me from discussing his proposals in detail, I will note that his proposals apparently only work for affirmative adjectives, and the above demonstration of the defeasibility of such entailment relations could be problematic for his theory.



I refer the reader to Vangsnes 1999 (from which this analysis has been taken and simplified) for an account of many noun phrase word orders in a wide variety of Scandinavian dialects. Crucially, this account relies on the fact that adjectives, as heads, count as interveners for head movement. It is unclear how such data could be elegantly captured on other assumptions concerning the phrase structural status of adjectives.

5 Conclusion

This paper has shown the advantages of an analysis of attributive adjectives as second-order identity-typed functions taking nominal arguments. The contrast between this analysis and the evidence that predicative adjectives behave as properties motivated the proposal of a type-shifting operator JOIN, deriving the attributive use of adjectives from the predicative use. It was then argued that attributive adjectives subcategorise for arguments of given syntactic categories, as well as of given semantic types. In particular, an attributive adjective subcategorises for a nominal sister. This suggests an analysis of the attributive adjective as the head of the Adj-N unit. A sketch of how this position might be accommodated within a model of lexical categories following Baker 2003 and Williams' (1994) theory of relativised heads was offered, before, finally, some further empirical consequences of the analysis of attributive adjectives as heads taking nominal complements were demonstrated.

References

- Abney, S. (1987). *The English Noun Phrase in its Sentential Aspect*. Ph. D. thesis, MIT, Cambridge, MA.
- Alexiadou, A. (2001). 'Adjective Syntax and Noun Raising: word order asymmetries in the DP as the result of adjective distribution'. *Studia linguistica* 55, 217–248.
- Baker, M. (2003). *Lexical Categories: verbs, nouns and adjectives*. Cambridge: Cambridge University Press.
- Bernstein, J. (1993). *Topics in the Syntax of Nominal Structure across Romance*. Ph. D. thesis, City University, New York.

- Chierchia, G. & R. Turner (1988). 'Semantics and Property Theory'. *Linguistics and Philosophy* 11, 261–302.
- Chomsky, N. (2000). Minimalist Inquiries: the framework. In R. Martin, D. Michaels, & J. Uriagereka (Eds.), *Step by Step: essays in minimalist syntax in honor of Howard Lasnik*, pp. 89–155. Cambridge, MA: MIT Press.
- Cinque, G. (1994). 'On the Evidence for Partial N Movement in the Romance DP'. In G. Cinque, J. Koster, J.-Y. Pollock, L. Rizzi, & R. Zanuttini (Eds.), *Paths Towards Universal Grammar*, pp. 85–110. Georgetown: Georgetown University Press.
- Grimshaw, J. (2003). Extended Projection. ms., Rutgers University.
- Heim, I. & A. Kratzer (1998). *Semantics in Generative Grammar*. Oxford: Blackwell.
- Higginbotham, J. (1985). 'On Semantics'. *Linguistic Inquiry* 16, 547–593.
- Kamp, J. A. W. (1975). 'Two Theories about Adjectives'. In E. Keenan (Ed.), *Formal Semantics of Natural Language*, pp. 123–155. Cambridge: Cambridge University Press.
- Kayne, R. (1994). *The Antisymmetry of Syntax*. Cambridge, MA: MIT Press.
- Lewis, D. (1972). 'General Semantics'. In D. Davidson & G. Harman (Eds.), *Semantics of Natural Language*, pp. 169–218. Dordrecht: D. Reidel.
- Montague, R. (1974). *Formal Philosophy*. New Haven: Yale University Press.
- Partee, B. (1987). 'Noun Phrase Interpretation and Type-shifting Principles'. In J. Groenendijk, D. de Jongh, & M. Stokhof (Eds.), *Studies in Discourse Representation Theory and the Theory of Generalised Quantifiers*, pp. 115–143. Dordrecht: Foris.
- Schwarzschild, R. (1999). GIVENNESS, AVOIDF and other Constraints on the Placement of Accent. *Natural Language Semantics* 7, 141–177.
- Selkirk, E. (1995). Sentence Prosody: intonation, stress and phrasing. In J. Goldsmith (Ed.), *The Handbook of Phonological Theory*, pp. 550–569. Oxford: Blackwell.
- Vangsnes, Ø. (1999). *The Identification of Functional Architecture*. Ph. D. thesis, University of Bergen.
- Wagner, M. (2005). Givenness and Locality. Paper presented at University College, London, 26/1/2005.
- Williams, E. (1994). *Thematic Structure in Syntax*. Cambridge, MA: MIT Press.