Reflexives and TAG Semantics

Robert Frank
Department of Cognitive Science
Johns Hopkins University
3400 N. Charles Street
Baltimore, MD 20910, USA
rfrank@jhu.edu

Abstract

Nesson and Shieber (2006) argue that the synchronous TAG (STAG) formalism provides an empirically adequate, yet formally restrictive solution to the problem of associating semantic interpretations with TAG derivations. In this paper, I further explore this approach, focusing on the semantics of reflexives. I find that STAG indeed permits a simple analysis of core cases of reflexives. This analysis does not, however, easily extend to contexts in which the reflexive and its antecedent are arguments of distinct elementary trees. I consider three possible extensions to the analysis which remedy these difficulties.

1 Introduction

The TAG community has recently witnessed an explosion of research into the problem of assigning semantic interpretations to TAG derivations. One line of work, beginning with Shieber and Schabes (1990), uses the synchronous TAG (STAG) formalism to build syntactic and logical form representations in parallel. The second type of proposal, put forward originally by Kallmeyer and Joshi (2003) and refined and extended in Kallmeyer and Romero (2008), exploits a unification operation defined over semantic feature structures associated with elementary trees to produce a Minimal Recursion Semantics representation. Nesson and Shieber (2006) argue that because the STAG proposal makes use of no additional machinery beyond the TAG formalism itself, it provides a more restrictive solution to the problem of semantic interpretation. To the degree that STAG is adequate to the task, one should then prefer it, as it holds out the possibility for providing explanatory accounts of semantic phenomena, much as TAG's restrictiveness has been shown to yield explanatory accounts of syntactic phenomena (Kroch, 1987; Frank, 2002). In this paper, I explore the question of STAG's adequacy, focusing on the phenomenon of reflexive interpretation.

As is well-known, reflexives are referentially dependent elements which are interpreted through their relation with a syntactically local antecedent. This syntactic sensitivity has led to analyses of the distribution of reflexives in terms of a syntactic constraint on the establishment of a syntactic correlate of the antecedent-reflexive relation (i.e., indexations), most famously the Binding Theory of Chomsky (1981). An alternative approach, explored by Partee and Bach (1984) and in much work since, assumes that anaphoric dependencies are instead established during the process of computing a semantic interpretation. The essential idea in such treatments is that reflexives are higher order functions over (transitive) predicates, decreasing the arity of the predicate by one and identifying the semantic value of two of the predicate's arguments via lambda abstraction.

(1) $[\![\text{himself}]\!] = \lambda P_{\langle e,\langle e,t\rangle\rangle} \lambda x.P(x,x)$

2 An STAG analysis

We can mimic this semantic treatment of reflexives in STAG by using the elementary tree set given in Figure 1a. The multicomponent tree set includes, on the syntactic side, trees corresponding to the reflexive and its antecedent, constrained to stand in a syntactic c-command relation. This is identical to the assumptions of Ryant and Schef-

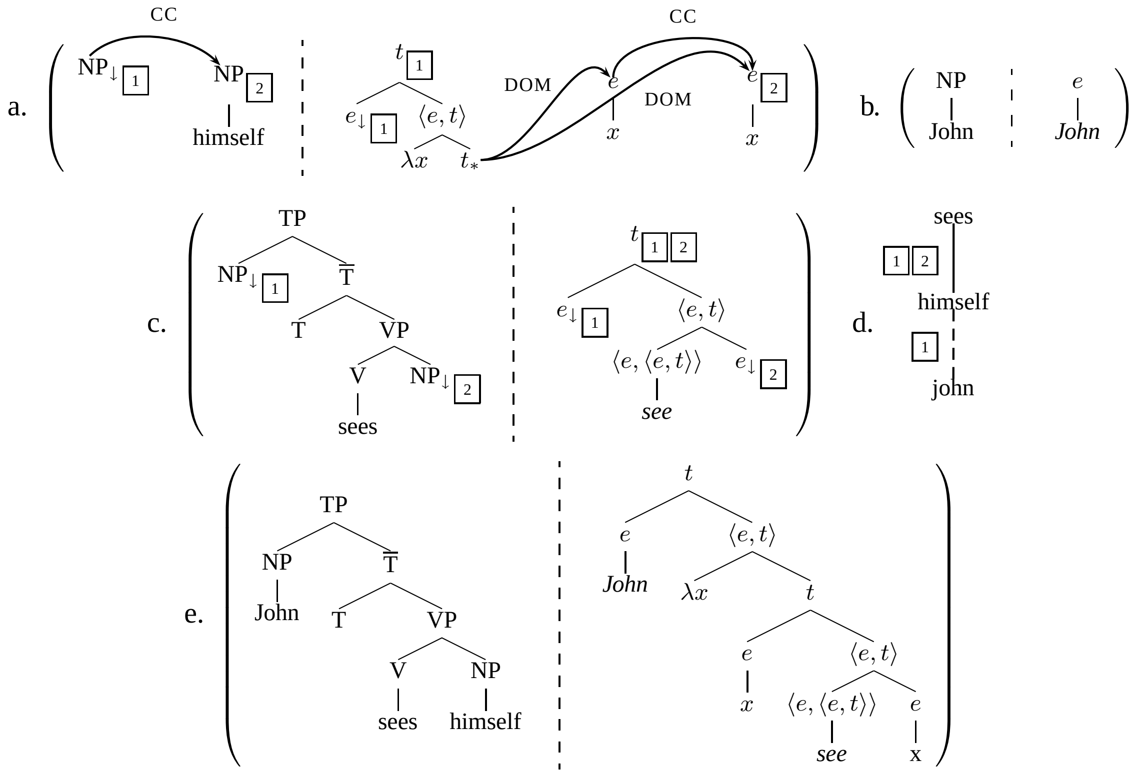

Figure 1: Elementary trees for (a) reflexives, (b) type e nominals, and (c) transitive predicates; (d) is the derivation tree for *John sees himself*, and (e) is the derived tree.

fler (2006) who propose an analysis of reflexives in the unification-based framework. On the semantic side, the tree set includes one tree that represents an instance of function application and lambda abstraction and two other trees, each an instance of the variable over which abstraction has taken place, constrained to stand in the relevant c-command and dominance relations. This tree set and the other trees depicted in Figure 1b and c can be employed in the tree-local derivation represented by the derivation tree in Figure 1d to produce the derived trees in Figure 1e. In this derivation, the antecedent first substitutes into the reflexive elementary tree, and the result then composes into the verbally headed elementary tree.

This type of derivation works equally well with quantified subjects. For such a case, I assume the semantic representation of the quantifier familiar from other STAG-based semantics work, shown in Figure 2a. To generate such an example, we follow the derivation depicted in Figure 2b. First, the quantifier combines first with reflexive: on the syntax side the NP tree representing the quantifier substitutes into the degenerate NP tree from the reflexive's tree set, while on the semantic side, the t-recursive auxiliary tree from the quantifier interpretation adjoins to the root of the t-recursive auxiliary from the reflexive interpretation and the e-rooted variable substitutes into the substitution slot in the same tree, thereby satisfying tree-locality. (Another derivation with adjoining to the foot rather than to the root would also satisfy tree-locality, but would violate the dominance restriction imposed by the quantifier tree set that ensures variable binding.) The resulting multi-component set is then combined with the verbally-headed elementary tree, as in the previous derivation, to produce the derived trees in Figure 2c.

This analysis extends to examples with reflexives embedded in non-quantificational picture-NPs:

(2) John bought the picture of himself.

To derive such a case, we need only adjoin the pair of trees depicted in Figure 3a, representing the head of the picture-NP, to the root of the reflexive-headed NP tree and its semantic analog. The derivation for (2) then continues just as in Figure 1.

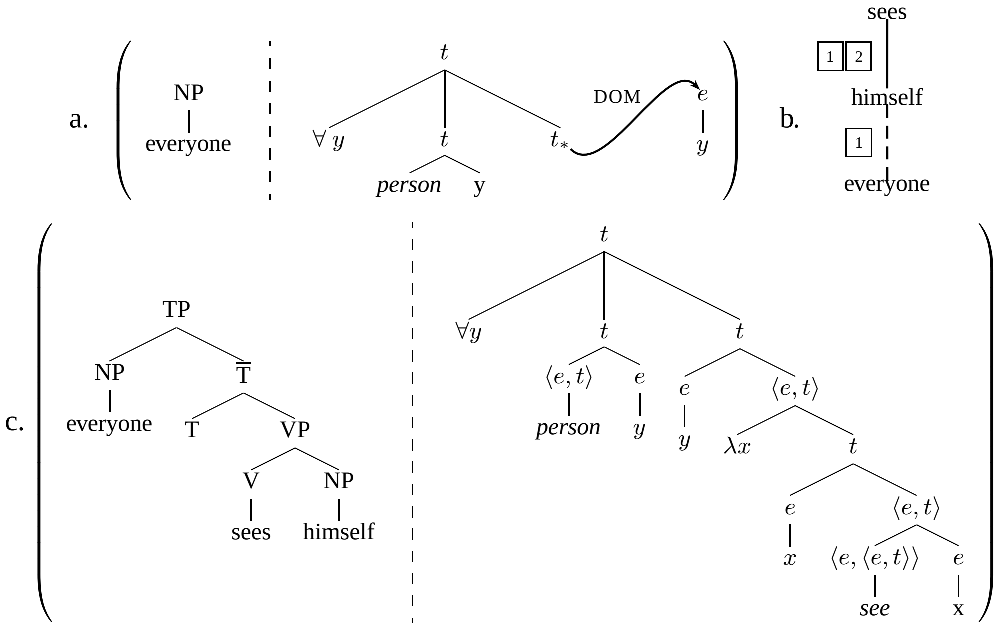

Figure 2: Derivation for quantifier-bound reflexives in *everyone sees himself*: (a) quantifier elementary trees, (b) derivation tree, (c) derived trees

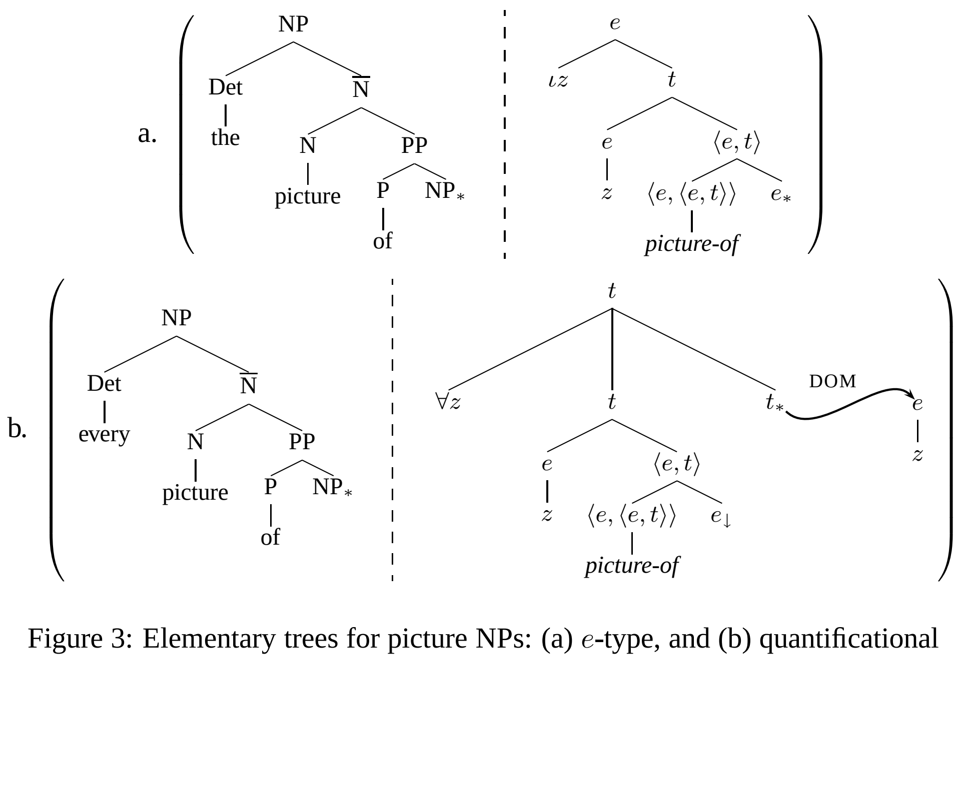

Figure 3: Elementary trees for picture NPs: (a) e-type, and (b) quantificational

One might object that the synchronous elementary tree set I have associated with the reflexive in Figure 1a is incompatible with the TAG version of the theta criterion (Frank, 2002), according to which all non-projected nodes need to be licensed via some predicate-argument relation and all arguments must be represented as a non-projected node. Under the reasonable assumption that this constraint applies to syntactic and semantic elementary tree sets alike, just as it applies to individual elementary trees, it is not clear how the syntactic and semantic elementary tree sets that represent the reflexive can be both well-formed: the syntactic set includes a single non-projected node (corresponding to the NP that is the antecedent of the reflexive), while the semantic set includes two (the substitution node of type e and the foot node of type t). This distinction suggests that the TAG theta criterion should more properly be understood as the reflection of a more general constraint on the expression of grammatical dependencies in elementary trees or tree sets, whether syntactic or semantic. The relevant dependency in the case of the syntactic representation of the reflexive is the relation established with its antecedent. The semantic dependencies, in contrast, are those that can be read off of the semantic interpretation in (1): the reflexive denotes a relation between an individual and a predicate and it is these that are realized as non-projected argument slots in this tree set.[1]

Divergences between the syntactic and semantic dependents will be found quite widely in functional elements whose denotations are taken to be higher order functions. Such elements will include quantifiers, reflexives, measure heads, relative pronouns and wh-phrases. In the case of lexical predicates, the syntactic and semantic dependents will tend to be better aligned though even here there may be divergences. The landing site of a wh-phrase that has undergone wh-movement to a higher clause might be thought of as a kind of syntactic dependent to the higher clause, and indeed the treatments of 'long' movement by Frank and Kroch (1995) and Frank (2002) in which the wh-phrase is substituted into such a position can be thought of as adopting this perspective. Similarly, one might be tempted to represent in the syntactic side of a synchronous grammar the case dependency between a finite raising verb and its subject or between an ECM verb and the subject of its complement clause (cf. Carroll et al. (2000)), while maintaining the standard set of dependencies on the semantic side. Pursuing this line of analysis raises a host of issues that lie beyond the scope of the current work.

This STAG analysis has a couple of significant advantages over compositional treatments using meanings like the one in (1). First of all, it avoids the need to multiply interpretations for the reflexive when it occurs as the dative argument of ditransitive predicates or with different antecedents. To derive the interpretive possibilities in (3) and (4), the different interpretations shown below each example must be assigned to the reflexive.

(3) Mary showed John himself in the mirror.
$[\![\text{himself}]\!] = \lambda P_{\langle e,\langle e,\langle e,t\rangle\rangle\rangle}\lambda x\lambda y.P(y,x,x)$

(4) John showed Mary himself in the mirror.
$[\![\text{himself}]\!] = \lambda P_{\langle e,\langle e,\langle e,t\rangle\rangle\rangle}\lambda x\lambda y.P(y,x,y)$

Under the STAG analysis, both of these interpretations can be derived from the single reflexive tree set in Figure 1. The difference between the different binding possibilities depends on the locus of substitution for the degenerate NP elementary tree from the reflexive tree set, and correspondingly the locus of substitution for the lambda-bound variable, whether into the patient or goal argument slots.

Secondly, the syntactic locality of the reflexive-antecedent relation derives not from a stipulation on semantics of the reflexive, but rather from the local nature of the TAG derivation. In contrast, using a semantic calculus using a denotation for the reflexive such as (1) as well as the operation of function composition, one could compute an interpretation of type $\langle e,\langle e,t\rangle\rangle$ for the word sequence *thinks that Mary admires*. As Szabolcsi (1987) notes, such a unit could then be combined with the reflexive and subject NP to yield a long-distance interpretation for the reflexive, an inter-

[1] Note that the predicate that is an argument of the reflexive is of type t in the TAG tree set, as opposed to the type $\langle e,\langle e,t\rangle\rangle$ in (1). This is a result of the flexibility afforded by multi-component composition as compared to function application in the more standard semantic calculus. It would be interesting to see whether such felxibility would allow us to restrict the types of all non-projected argument nodes to base (as opposed to function) types. Such a restriction would, if it can be maintained, impose substantial restrictions on possible interpretations for lexical elements.

pretation that must be blocked via some additional stipulation.[2]

3 Moving beyond clausemates

Attractive as this analysis is, it has two shortcomings if it is to serve as a demonstration of the viability of STAG semantics in this domain. First of all, so long as the strictures of tree-local (or even set-local) MCTAG are maintained, the analysis cannot be extended to cases of reflexives embedded in quantificational picture NPs like the following:

(5) John bought every picture of himself.

The derivation of such an example will not involve the e-type tree set for the picture NP in Figure 3a, but must instead use the quantificational tree set in Figure 3b. This tree set cannot however be combined with the reflexive tree set in a tree-local or even set local fashion. On the one hand, the reflexive trees cannot both adjoin or substitute into the *picture* tree on the syntactic side, since the latter provides no position for the antecedent. The reverse combination can proceed on the syntactic side of the derivation, where the picture NP tree adjoins to the root of the reflexive-headed NP. However, on the semantic side this do not work out, since this NP is linked only to the e-type variable tree, which can host either adjoining of the t-recursive auxiliary tree or substitution of type e variable. Note that even if the root of the reflexive NP were linked to the root of the t-recursive auxiliary tree, allowing for the scope to be established, the derivation would still fail because of the absence of a slot for substitution of the variable introduced by the picture-NP tree. Both of these failures arise from the same source, neither the reflexive nor the picture-NP tree include structural representation of the predicate one of whose semantic arguments needs to be quantified over, and one of whose syntactic arguments needs to serve as the antecedent for the reflexive.

We need not necessarily despair at this aspect of our analysis. On the basis of examples like (6) and (7), Pollard and Sag (1992) and Reinhart and Reuland (1993) have argued that reflexives inside of picture-NPs are in fact exempt from the usual syntactically-defined locality conditions on reflexive interpretation.

(6) Bill_i finally realized that if The Times was going to print [that picture of himself_i with Gorbachev] in the Sunday edition, there might be some backlash.

(7) Lucie_i said that (you agreed that) a picture of herself_i would be nice on that wall.

These authors argue that the interpretation of reflexives in picture-NP contexts is determined by pragmatically defined conditions.

Even if we suppose that this is correct, this is not enough to avoid difficulties entirely, as there is another pair of constructions that leads to problems: raising and ECM.

(8) John seems to himself to be the best candidate.

(9) John considers himself to be the best candidate.

Let us turn first to raising. Under the usual TAG derivation of a raising sentence like (8), the raising verb is represented by an auxiliary tree that lacks a position for a subject. This lack of a subject position immediately causes a problem when we attempt to combine the reflexive tree set with the raising auxiliary of which it is an argument, as there is no position that serve as the attachment site for the degenerate NP tree (see Figure 1a), and therefore we cannot retain either tree- or set-locality.[3]

There are a number of possible lines of analysis we might pursue here. I will outline each briefly, but space prevents me from deciding among them. The first involves a rethinking of the TAG syntactic analysis of raising, along the lines envisioned in the previous section, so that the syntactic representation of the raising verb's elementary tree would indeed include a representation of its syntactic dependent, the subject. This would permit the incor-

[2] Szabolcsi (1987) argues that the lack of locality built into the reflexive's semantics is desirable in order to deal with cases of long-distance anaphors (Koster and Reuland, 1991). Clearly, my current proposal does not extend to such cases, and must treat them via a different mechanism.

[3] For such reasons, Ryant and Scheffler (2006) in their analysis of reflexives exploit the flexible composition operation, thereby losing much of the constraint that the TAG formalism imposes on derivations. (Kallmeyer and Romero, 2007) demonstrate that this problem can be avoided by taking the antecedent component of the syntactic representation of the reflexive to adjoin to VP rather than NP. This move requires however the introduction of a regimen of feature passing of antecedent features.

poration of both halves of the reflexive's (syntactic) tree set into this the raising predicate's tree set. I will leave this option unexplored, because of the broad implications it would have on the treatment of locality in raising constructions more generally.

A second option involves taking the combination of the raising predicate and the reflexive experiencer to be the result of a lexical process, so that this combination was represented via a single elementary tree set. As seen in Figure 4, this tree set would have two components in the syntactic half to incorporate the representation of the reflexive's antecedent. This tree set could be adjoined into an infinitival clause to produce the appropriate syntax and interpretation. Such a lexicalist analysis of reflexives could in fact be applied to the monoclausal and picture NP cases discussed earlier, as well as to ECM.[4]

Unlike the first two possibilities which retain the same two part elementary tree set for the reflexive, a third analytic option alters this assumption. Specifically, this analysis adopts the considerably simpler view of reflexive syntax and semantics represented in Figure 5a, according to which the syntactic representation of of a reflexive is a single NP elementary tree, and the interpretation is variable of type e. When this reflexive is substituted into a raising auxiliary tree, the identity of the variable with which it is associated is percolated to the root of the raising auxiliary, as in the tree in Figure 5b. To accomplish the binding of this variable, I assume that the syntax and logical form associated with a simple clause are both somewhat more complex than we have been assuming, but in a manner that has independent motivation (see Figure 5c). On the syntax side, I take the subject to be generated within VP and raised to its surface position. On the semantics side, this structural assumption translates into a lambda-bound variable that is saturated by the surface subject (Heim and Kratzer, 1998). Because the VAR feature ensures the identity of the reflexive variable and the lambda bound variable, the reflexive variable will be bound once the raising auxiliary adjoins into the infinitival clause. A similar analysis will work for the ECM case as well, using the pair of trees in Figure 6 as the representation for the ECM predicate.

The VAR feature, if it is to fit with the feature system of TAG as usually understood, can take only one of a finite set of values.[5] In fact, such a bound on the number of distinct variables that can be present in an STAG-derived logical form is already imposed upon us by the fact that these logical forms are constructed from a TAG, which by definition may contain only a finite set of elementary trees. The restriction to a bounded number of distinct variables does not, of course, rule out the generation of sentences with unbounded complexity, so long as variables can be "reused". Because of the bounded nature of reflexive binding, the restriction will not cause any difficulties as the domain over which a reflexive can be bound is limited. What the restriction does rule out is, for instance, sentences with interactions between unboundedly many quantifiers and variables bound by them. As the number of such quantifier-pronoun pairings increases, such sentences become ever more difficult to comprehend, and it is therefore a rather thorny theoretical question as to whether such examples ought to be generated by the grammar (Joshi et al., 2000).

References

Carroll, John, Nicolas Nicolov, Olga Shaumyan, Martine Smets, and David Weir. 2000. Engineering a wide-coverage lexicalized grammar. In *Proceedings of the 5th International Workshop on Tree Adjoining Grammars and Related Frameworks*, pages 55–60, Paris.

Chomsky, Noam. 1981. *Lectures on Government and Binding*. Foris, Dordrecht.

Frank, Robert and Anthony Kroch. 1995. Generalized

[4] The ECM cases raise a problem as they might allow locality to be circumvented by repeated adjoining to the root of the syntactic tree representing the infinitival clause, hosting the reflexive. I leave this issue open for future work. Interestingly, such an issue does not arise in the case of raising, as the $\overline{\text{T}}$ recursion of the raising auxiliary prevents the introduction of an intervening antecedent even under repeated adjoining of auxiliary trees, without resort to intermediate traces. Interestingly, this analysis correctly predicts that intervening experiencers should not count as potential intervening binders.

(i) John$_i$ seems to Mary to appear to himself$_i$ to be the best candidate.

[5] This is a respect in which the feature unification exploited in work in TAG semantics is different from that that has traditionally been used in TAG syntax. And it is in this sense that the kind of feature unification I am exploiting is distinctly less powerful than that used by Kallmeyer and Romero (2007).

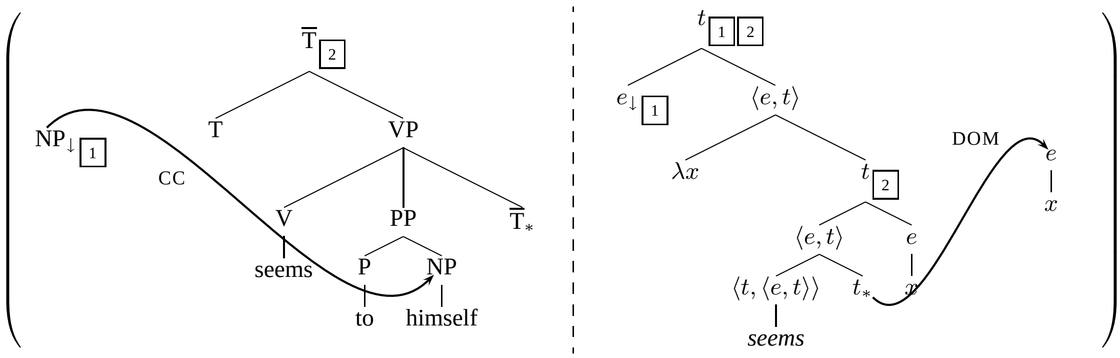

Figure 4: Elementary trees for composed reflexive-raising predicate

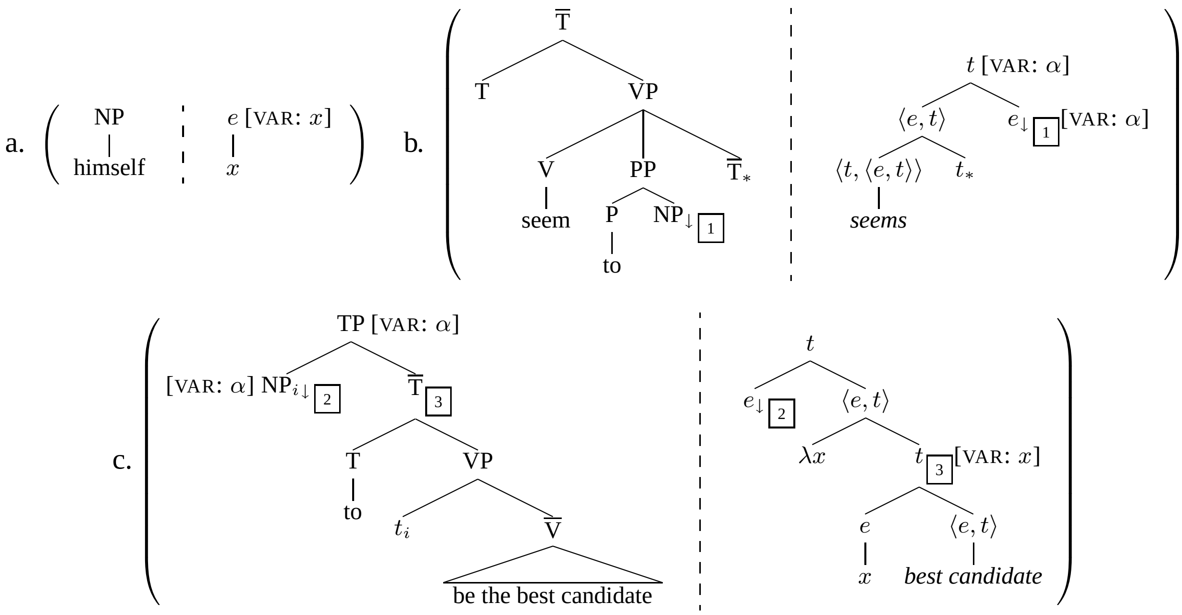

Figure 5: Elementary trees for raising derivation with VAR feature passing

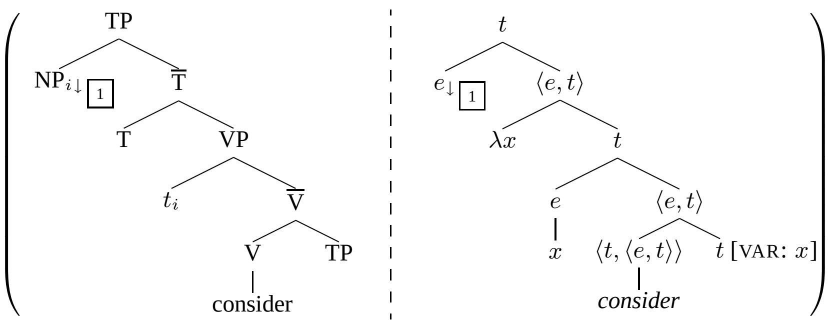

Figure 6: Elementary trees for ECM derivation with VAR feature passing

transformations and the theory of grammar. *Studia Linguistica*, 49(2):103–151.

Frank, Robert. 2002. *Phrase Structure Composition and Syntactic Dependencies*. MIT Press, Cambridge, MA.

Han, Chung-hye. 2007. Pied-piping in relative clauses: Syntax and compositional semantics using synchronous tree adjoining grammar. *Research on Language and Computation*, 5(4):457–479.

Heim, Irene and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Blackwell, Oxford.

Joshi, Aravind K., Tilman Becker, and Owen Rambow. 2000. Complexity of scrambling: A new twist on the competence-performance distinction. In Abeillé, Anne and Owen Rambow, editors, *Tree Adjoining Grammars: Formalisms, Linguistic Analysis and Processing*, pages 167–181. CSLI Publications, Stanford, CA.

Kallmeyer, Laura and Aravind K. Joshi. 2003. Factoring predicate argument and scope semantics: Underspecified semantics with LTAG. *Research on Language and Computation*, 1:3–58.

Kallmeyer, Laura and Maribel Romero. 2007. Reflexives and reciprocals in LTAG. In Geertzen, Jereon, Elias Thijsse, Harry Bunt, and Amanda Schiffrin, editors, *Proceedings of the Seventh International Workshop on Computational Semantics (IWCS-7)*, pages 271–282, Tilburg.

Kallmeyer, Laura and Maribel Romero. 2008. Scope and situation binding in LTAG using semantic unification. *Research on Language and Computation*, 6(1):3–52.

Koster, Jan and Eric Reuland, editors. 1991. *Long-Distance Anaphora*. Cambridge University Press, Cambridge.

Kroch, Anthony. 1987. Unbounded dependencies and subjacency in a tree adjoining grammar. In Manaster-Ramer, Alexis, editor, *The Mathematics of Language*, pages 143–172. John Benjamins, Amsterdam.

Nesson, Rebecca and Stuart M. Shieber. 2006. Simpler TAG semantics through synchronization. In Winter, Shuly, editor, *Proceedings of the 11th conference on Formal Grammar*, pages 129–142. CSLI Publications.

Partee, Barbara and Emmon Bach. 1984. Quantification, pronouns, and VP anaphora. In Groenendijk, Jeroen, Theo M.V. Janssen, and Martin Stockhof, editors, *Truth, Interpretation and Information*, pages 99–130. Foris Publications, Dordrecht.

Pollard, Carl and Ivan Sag. 1992. Anaphors in English and the scope of binding theory. *Linguistic Inquiry*, 23.

Reinhart, Tanya and Eric Reuland. 1993. Reflexivity. *Linguistic Inquiry*, 24:657–720.

Ryant, Neville and Tatjana Scheffler. 2006. Binding of anaphors in LTAG. In *Proceedings of the 8th International Workshop on Tree Adjoining Grammar and Related Formalisms (TAG+8)*, pages 65–72. Association for Computational Linguistics.

Shieber, Stuart and Yves Schabes. 1990. Synchronous tree adjoining grammars. In *Proceedings of the 13th International Conference on Computational Linguistics*, volume 3, pages 253–258, Helsinki.

Szabolcsi, Anna. 1987. Bound variables in syntax (are there any?). In Groenendijk, Jeroen, Martin Stockhof, and Fred Veltman, editors, *Proceedings of the 6th Amsterdam Colloquium*, pages 331–353, Amsterdam. Institute for Logic, Language and Information.